HEWLETT - PACKARD SOUTH QUEENSFERRY

Innovation in Silicon Glen 1965 - 2010

Hugh Walker

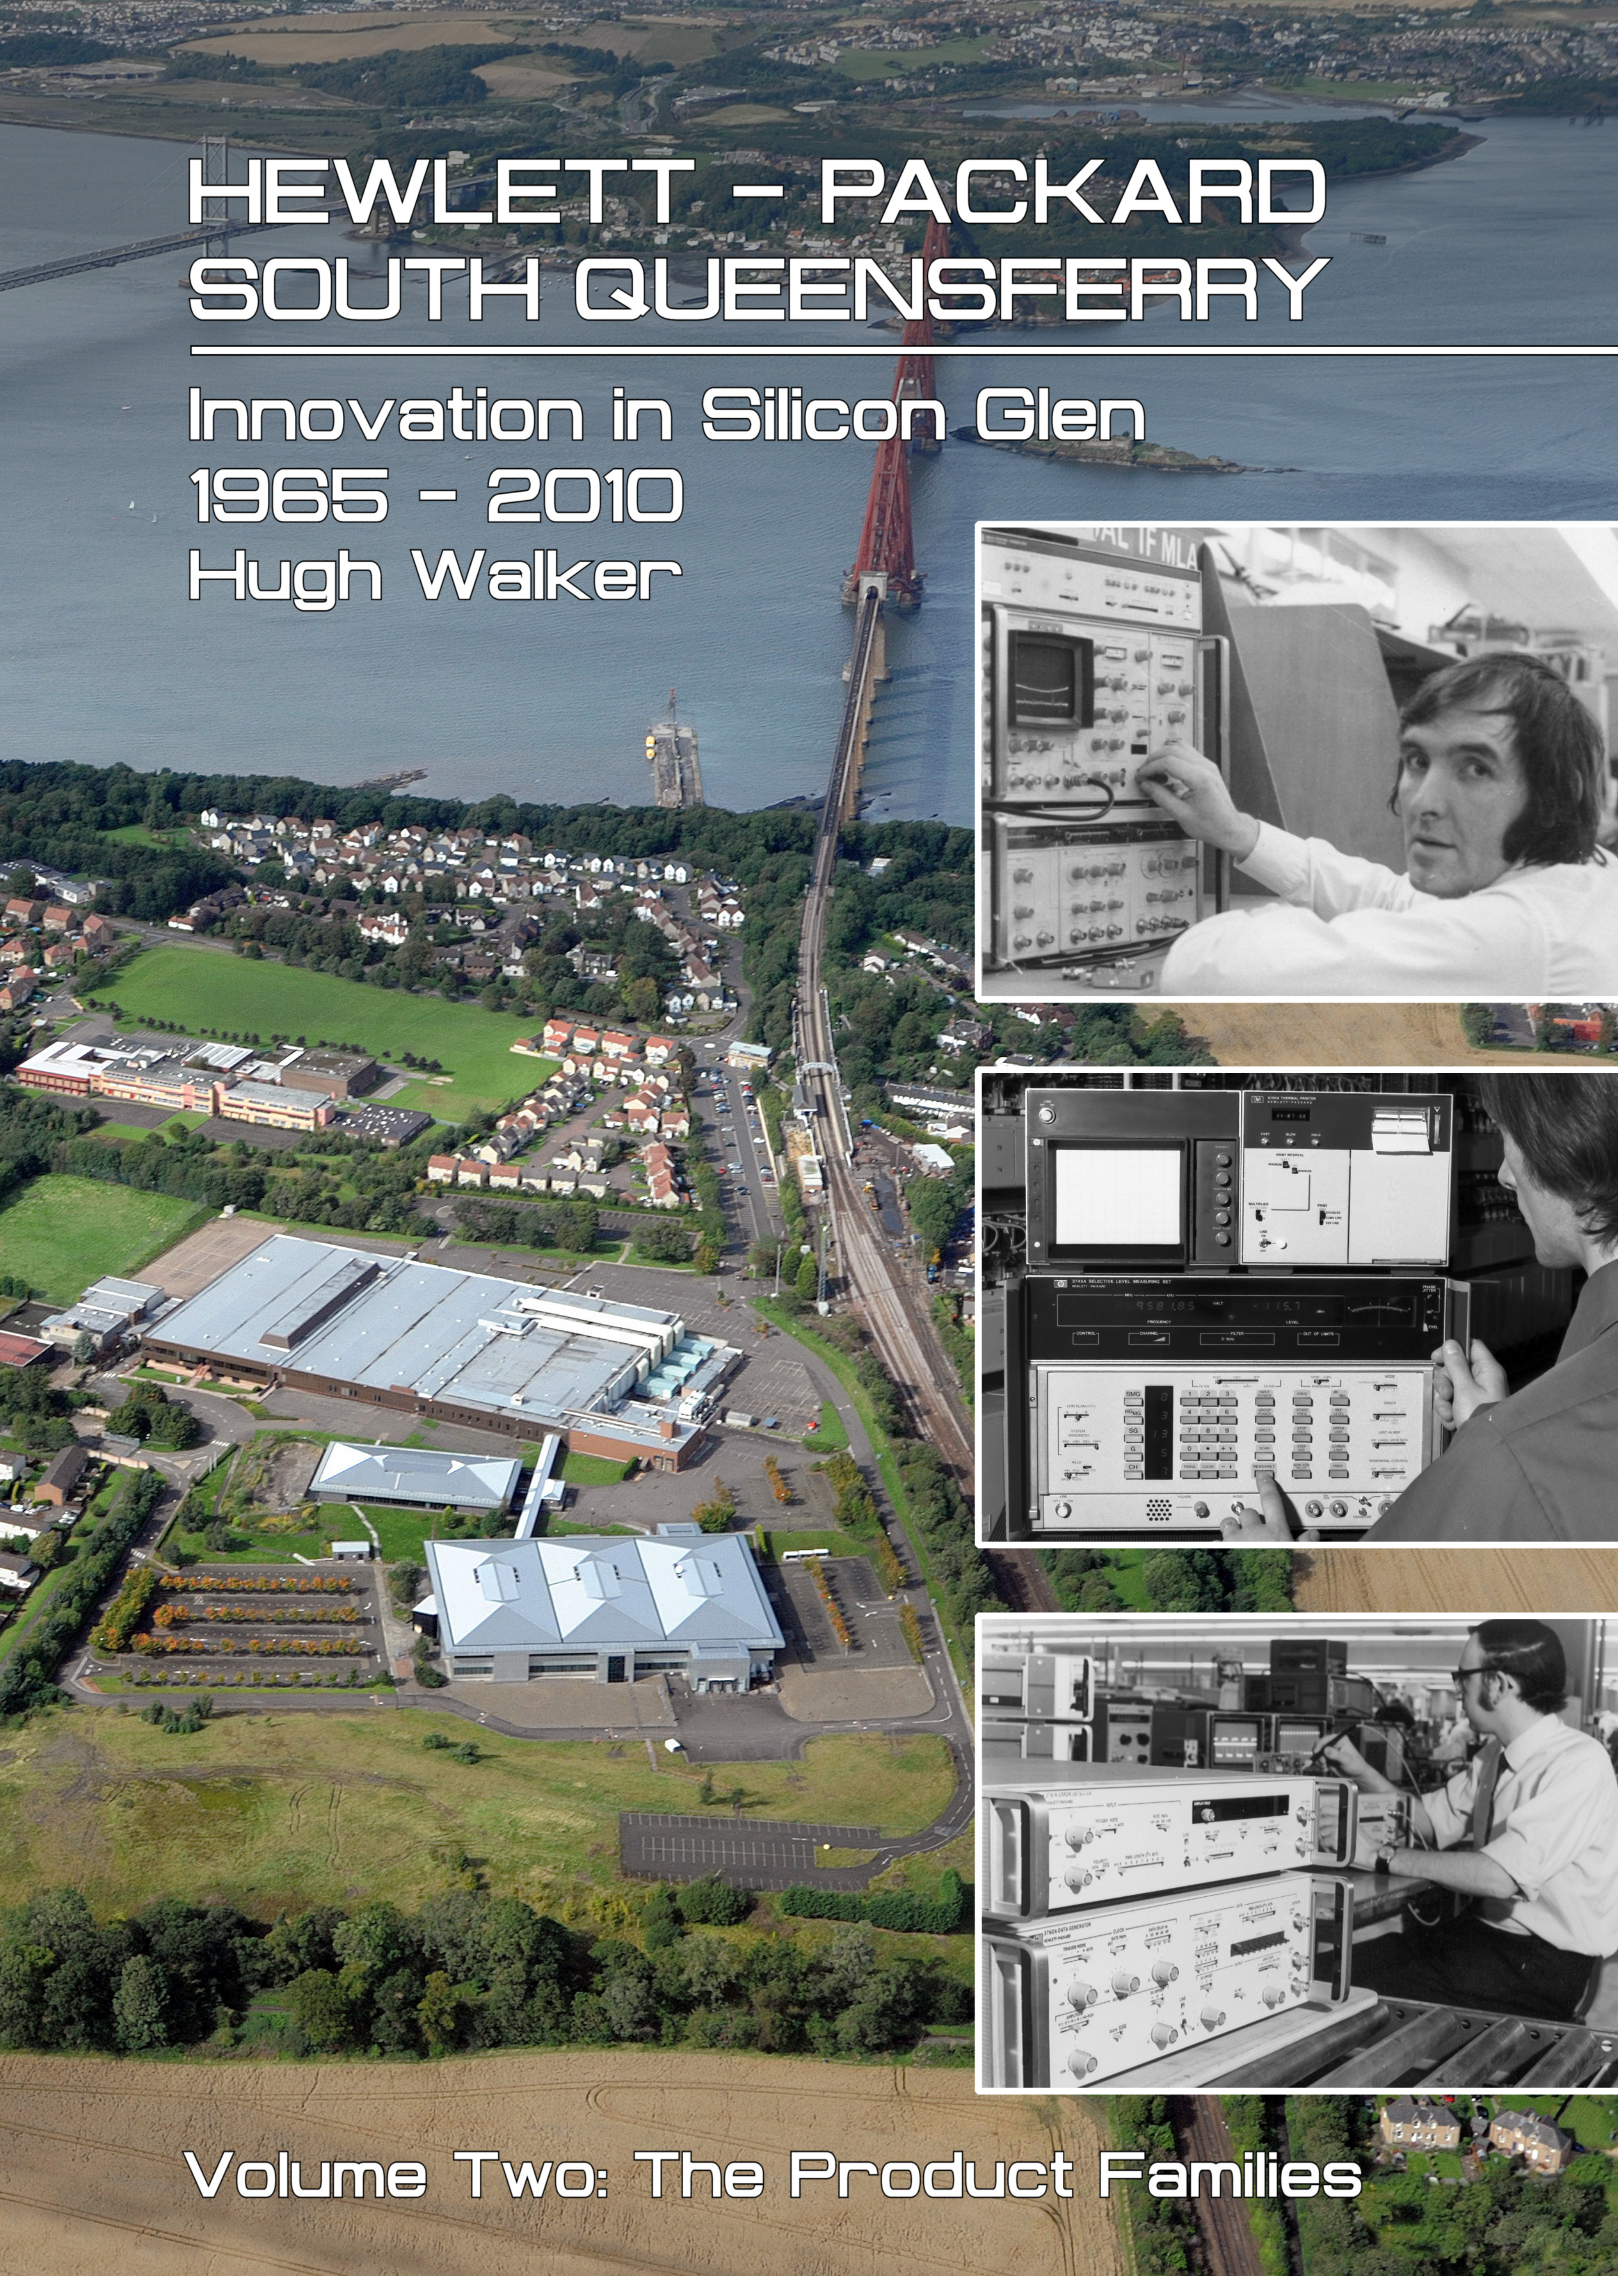

Volume Two: The Product Families

Product Families

As part of Hugh Walker's ten-year project in documenting the definitive history of HP in South Queensferry, this volume describes the development and marketing of the products designed by the talented engineers in QTD, South Queensferry.

With talented people working throughout the factory, it will be no surprise to read of the technological innovations achieved in the advancement of communications test and measurement.

This unique and fascinating story will be of interest to ex-employees and the families of HP factory and field staff who worked in, or closely with, HP in South Queensferry. It will also be of historical interest to academics and historians who study Silicon Glen.

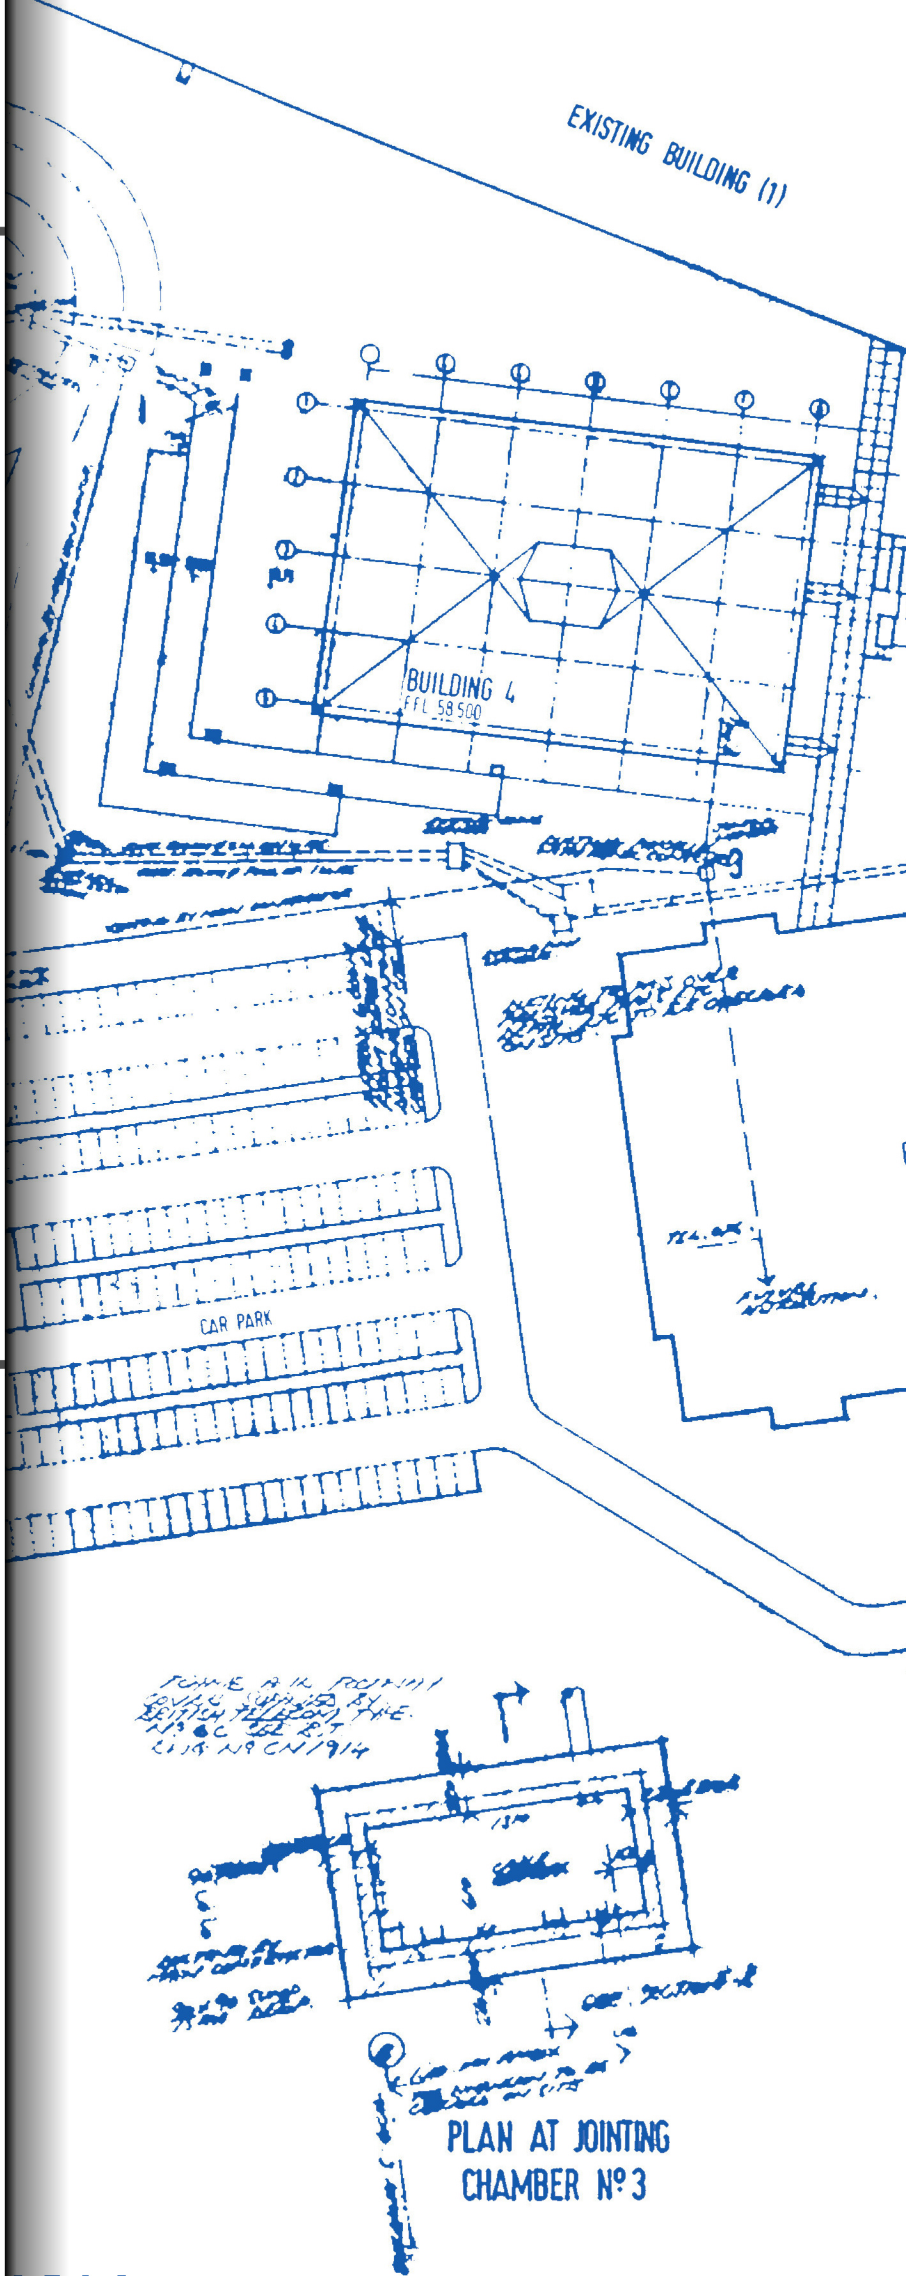

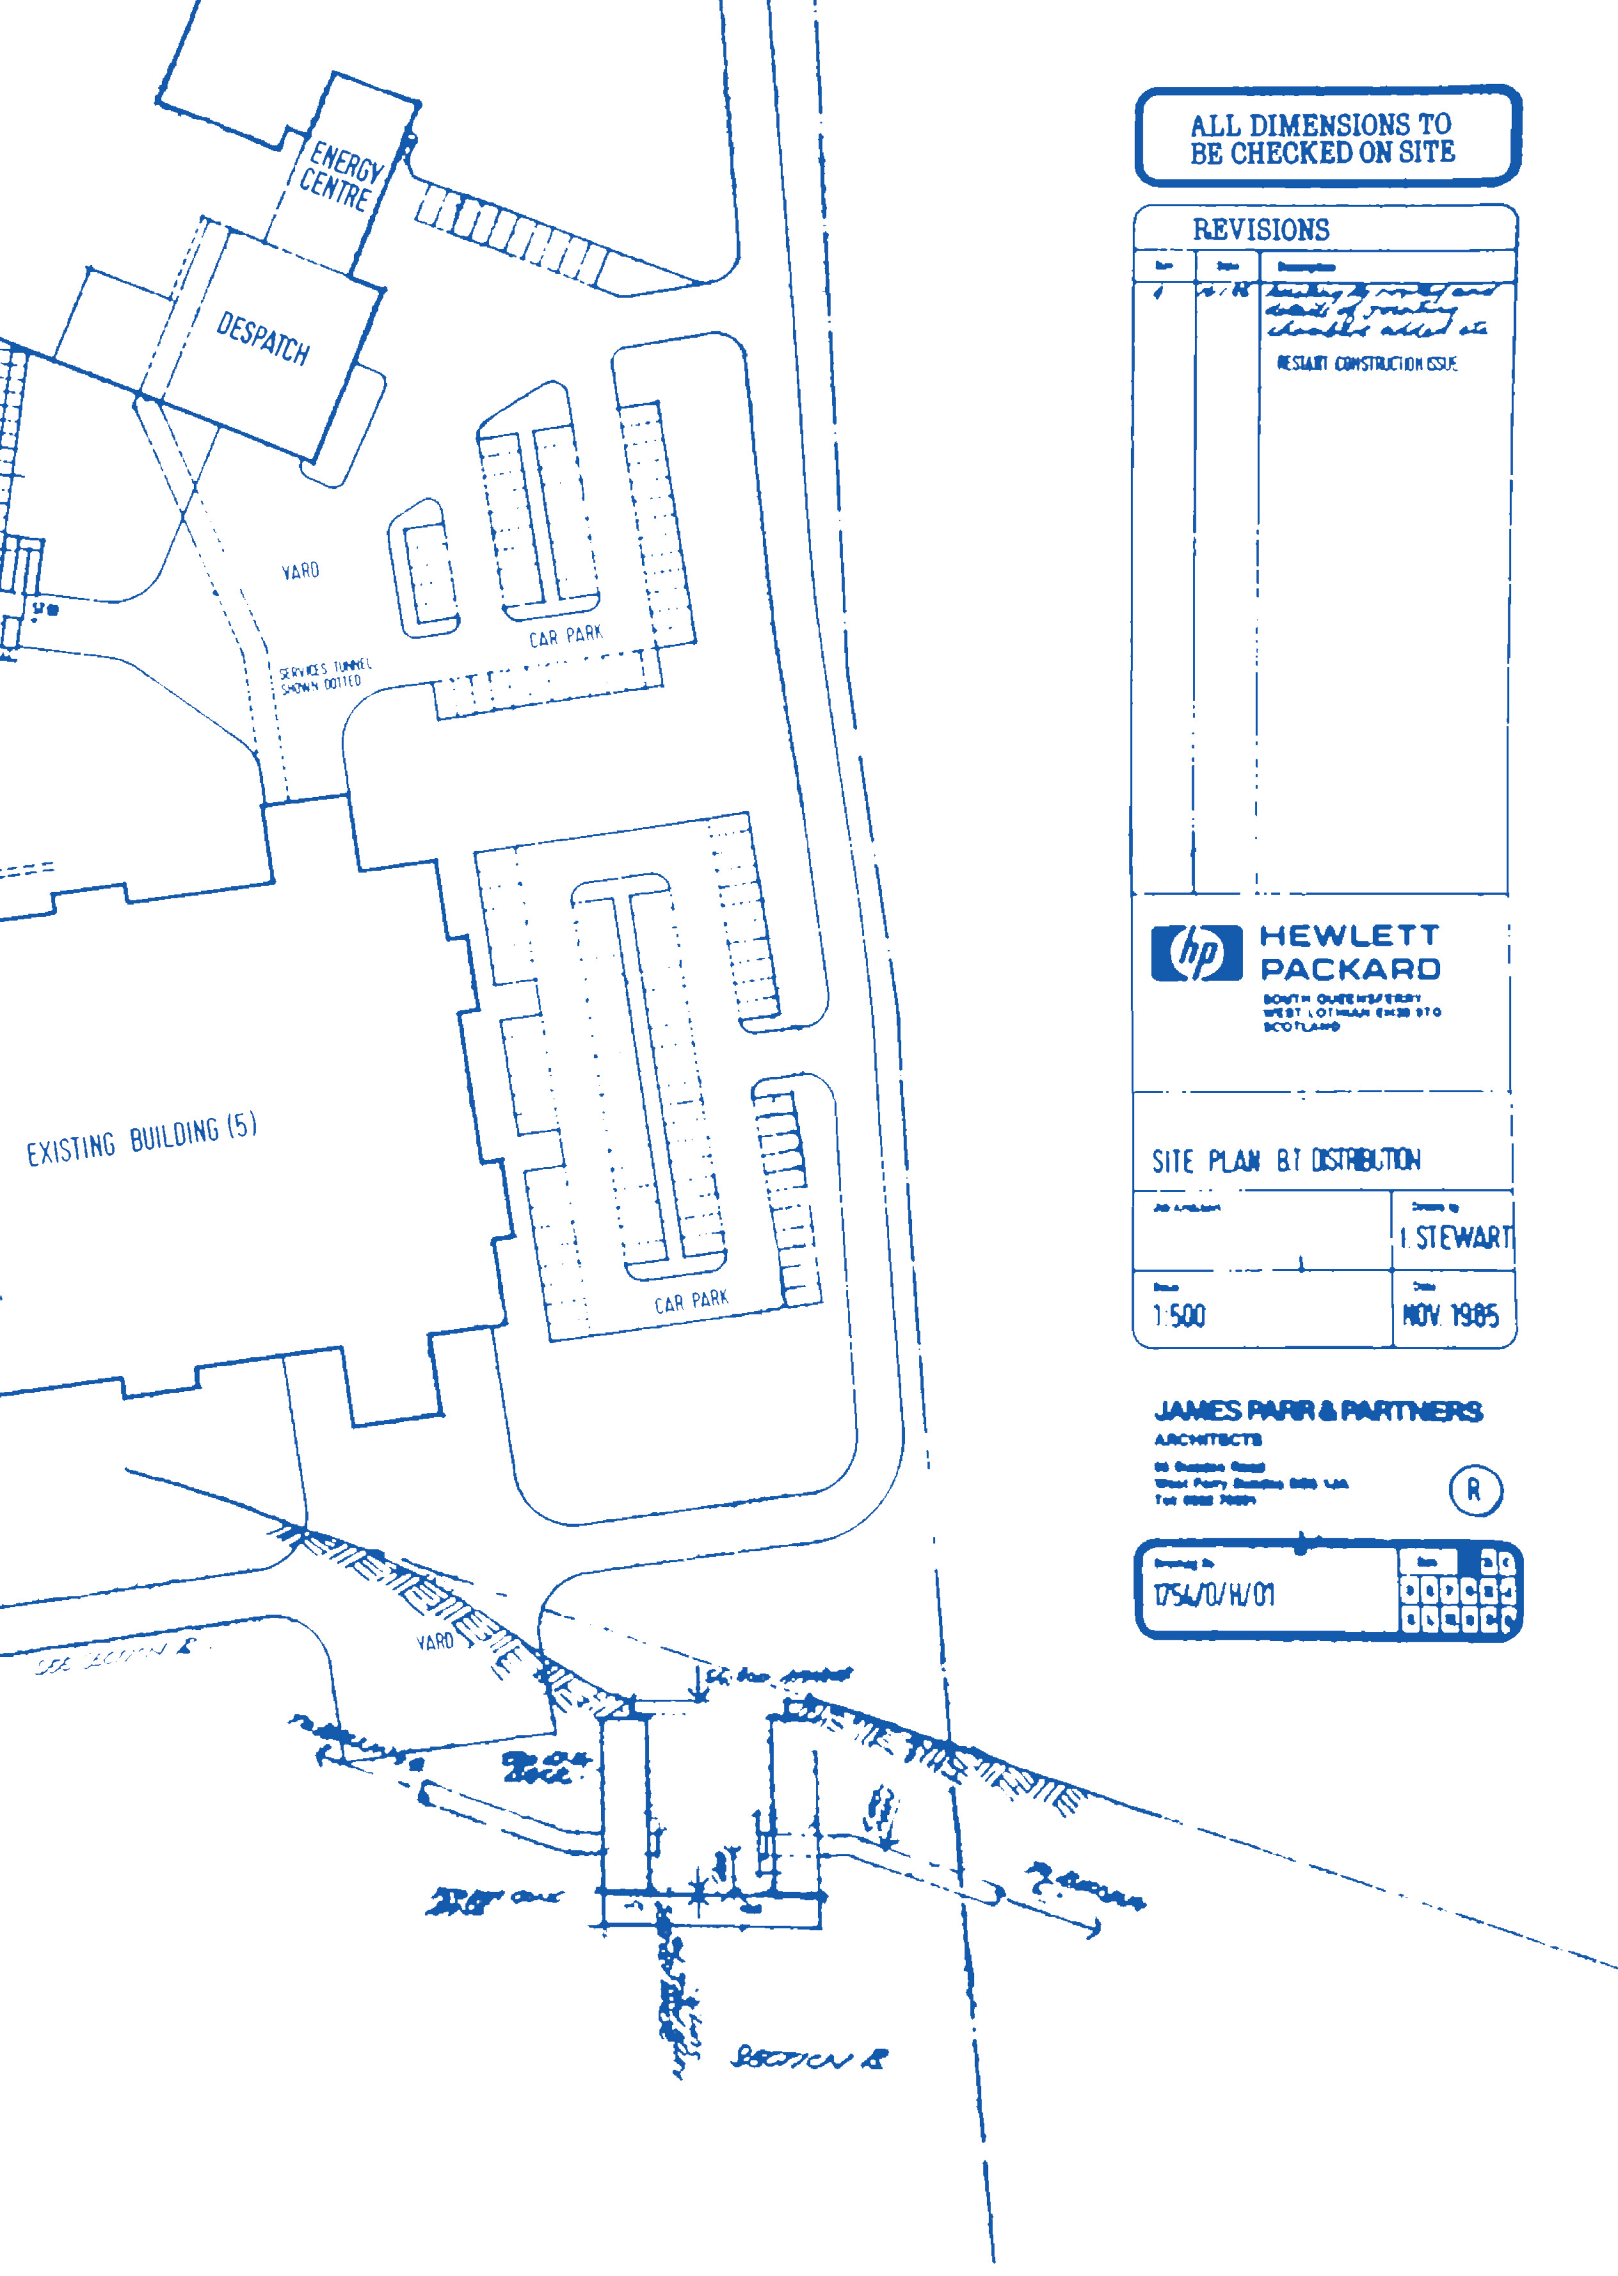
ENERGY CENTRE
DESPATCH
YARD
CAR PARK
SERVICES TUNNEL SHOWN DOTTED
EXISTING BUILDING (5)
CAR PARK
YARD
ALL DIMENSIONS TO BE CHECKED ON SITE
REVISIONS
RESTART CONSTRUCTION ISSUE
HEWLETT PACKARD
SITE PLAN BT DISTRIBUTION
I. STEWART
1:500
NOV. 1985
JAMES PARR & PARTNERS
ARCHITECTS
R
1754/0/H/01

HEWLETT - PACKARD
SOUTH QUEENSFERRY

Innovation in Silicon Glen
1965-2010

Cover Photographs:
Main image: Aerial View of the HP/Agilent Factory South Queensferry, 2011.
©2011 Tony Gorzkowski, Whitehouse Studios, Edinburgh.

Front insets: 3711/12 Microwave Link Analyser.
3745A Selective Level Measuring Set.
3760A Pattern Generator and 3761A Error Detector.

Rear insets: 37900D Signalling Test Set.
37718A OmniBER 718.
71612C HSBER.
37204A Bus Extender.

Front and rear flys: Artistic representation of architect’s site plan for Building 4 & 5.

This printed volume designed and produced by James Gentles.
Private printing by Swallowtail Print, Norwich.
First Printed in this form 2021 – 2a

Printed in Times New Roman, Calibri bold, and Probert.

While the printer is UK Forestry Stewardship Council (FSC) accredited and the paper chosen meets a similar EU regulation, we cannot use a FSC mark because there is no end-to-end control under the one standard.

HEWLETT - PACKARD
SOUTH QUEENSFERRY

Innovation in Silicon Glen
1965-2010

Volume Two:

The Product Families

by

Hugh Walker

Employee No: 66699

2014

This volume is dedicated to the many talented engineers at HP/Agilent South Queensferry in Scotland, who designed and marketed these instruments, and to all the manufacturing people and others who made and supported them.

About the Author

Hugh Walker was born in 1947 and brought up in Coventry. He has engineering degrees from Cambridge and Southampton Universities and joined HP at South Queensferry as a design engineer in October 1970. He had various design and management assignments in the R&D department, before moving to product marketing in 1982 where he had responsibility for analogue and digital transmission test. After a stint in the newly formed European Marketing Centre, he moved into marketing communications including PR in the late 1980s and spent the remainder of his career mainly in this activity. In the late 1990s, he also assisted the management team with market research. After taking early retirement at the end of 2001, he spent a few years as a freelance consultant in technical marketing for Scottish Enterprise and university commercialisation projects. Hugh is married to Janet who was the HP South Queensferry librarian in the 1970s. Their son, Mike, also worked at South Queensferry for 10 years as an IT contractor until the factory closed in 2010.

Contents

Appendices

Foreword

Finlay Mackenzie

It was in the early 1960s, while developing radar equipment at Ferranti in Edinburgh, that Peter Carmichael and I became aware of Hewlett-Packard (HP) measuring instruments, though we had no idea at that time that HP would soon play such a vital role in career development for both of us.

It is fair to say that Ferranti Edinburgh was the major player in Scottish Electronics at that time. They had large Research & Development (R&D) teams and good graduate training and apprenticeship schemes.
The expertise of these teams was built up over years and passed on to new graduates. They had low R&D engineer turnover. Here was a base of talent that would help spread the Electronics Industry in Scotland and help establish Silicon Glen.

In 1965 Peter Carmichael and I joined HP in Bedford, knowing they were going to set up manufacturing and R&D in South Queensferry, Scotland, the following year. Our first product, the Microwave Link Analyzer, won the Queen's Award for Industry and convinced Bill Hewlett and David Packard that our Product Development Charter at Queensferry should be to develop Telecommunications Test & Measuring Equipment dedicated to meet the needs of the world wide Telecom System Developers, Manufacturers and Maintenance Organisations. Other divisions in HP were developing general purpose test instrumentation for a range of markets, but not specifically for the Telecom Market.

In this early period, Queensferry also invented two non-telecom products, which later migrated to a product line elsewhere in HP: a Pseudo-Random Noise Generator and a Signal Correlator, used in mechanical analysis. Understanding this theory of the Pseudo-Random Binary Sequence (PRBS) was invaluable to our move into Digital Communications later on.

This emphasis on developing our own Product Line ensured HP's Inward Investment was not a so-called "Screwdriver Plant". As a result of the success of the product lines that Hugh has described in this Volume, HP South Queensferry established a great reputation in Silicon Glen. The Scottish Development Agency (SDA) used us very effectively to promote Inward Investment to Scotland. US Management Teams from Digital Equipment

Corporation (DEC) and Motorola would land their helicopters at South Queensferry because the SDA wanted them to see what HP had achieved there.

We were happy to see people leave to set up their own businesses. Indeed, we encouraged them. The first group to set up was Fortronics, followed by Livingston Precision, and there have been others more recently. When we ran out of space we added a new building so there was spare capacity and we allowed small start-up companies to rent space in what we called Incubation Units. Cash flow is usually a problem for start-ups, but many survived during the period of developing their own product, by getting some subcontract work from us. When they were established they moved to their own premises. These are some of the ways we contributed to the growth of Silicon Glen. Hopefully, because of our presence we were part of a "virtuous circle" that spread the R&D, Production and Product Marketing expertise of the Hewlett-Packard Company, which was "state of the art", throughout the UK. Some managers left to devote their full-time career to giving Seminars on such topics as "The Process of Management".

When we had 1000 employees at South Queensferry, Professor Donald Mackay of PIEDA Consulting estimated that the multiplier effect of our activity into the local and Business Community was over three times the above number. A decade after I retired there were over 2000 employees at Queensferry, creating thousands of additional jobs outside HP.

We were all enthusiastic and ready to tell anybody who would listen that we worked for a great Company. On work days, I always looked forward to going there and never missed a day through illness in over 25 years. There was such a feeling of "belonging and trust". For example, when business was slack in the 1980s, we all went on a nine-day fortnight at reduced pay so nobody was paid off. The spirit of the place was such that employees were happy that "everyone hurt a little so that no one hurt a lot". I also remember the owners of the Sealscraig Hotel in South Queensferry were astounded at how positive the production line girls were, on Friday nights, about their employer, when some other customers were typically running down their workplace. Going to work there was fun and never a grind.

From a business viewpoint, I soon realised on joining HP, that they put a lot of emphasis on Product Marketing. At HP South Queensferry we were spending 6% of our shipments on Product Marketing, excluding Field Selling costs. This really professional approach to the New Product Development Cycle is hopefully ingrained in the thought processes of the potential entrepreneurs who left HP/Agilent to set up their own businesses. Though HP South Queensferry is no longer there, perhaps Silicon Glen still benefits from the legacy derived from that "virtuous circle".

We were allowed to spend 10% of our Shipments on R&D. The success of our first product enabled us to go on the UK University Graduate Recruiting rounds, and bring engineers, whom we felt had a good interface with people, back for final selection by technical grilling at the factory. I met Hugh Walker for the first time during one of these sessions, and was really impressed with how he dealt with very difficult technical questions.

He accepted the R&D job offer and soon progressed to R&D section manager. Later, he decided that he fancied a job in Product Marketing, where he was welcomed with open arms. I was disappointed he was leaving R&D. He knew he had other talents of which we were unaware, and also a sense of history that led him to write two books on the Dunfermline linen industry in the early 1990s.

All who worked at HP South Queensferry should be grateful that Hugh has the talent and the desire to record the history of what was achieved at the factory for the last 35 years of the 20th century. This has been a major undertaking and Hugh had the foresight to realise that much of the documented information which still existed at South Queensferry was about to be shipped to the USA, probably never to surface again.

Hugh not only has a remarkable grasp of the basics of all the evolving Telecom Systems for which we were developing Test Products, but he also understands the detail at the system and circuit level of these Products, whose function and purpose he has so ably described. One measure of the internal strength of HP during the period of this history can be seen in the way friendships have continued, spanning decades and extending long after the products were developed. Hugh has been the catalyst who has brought out a purposeful enterprise from all these willing friends.

Dr. Finlay Mackenzie CBE, D.Eng, C.Eng, FIET
HP South Queensferry General Manager, 1982 to 1990.

Preface

Hugh Walker

In the spring of 2010, as the old HP South Queensferry facility entered its last few months, the site librarian, Mayumi Hepburn asked if I would come back to help look through the archives and photos stretching back 45 years and decide what should be done. My afternoon visit was timed to coincide with that of a librarian from Agilent Archives in Santa Clara, California. The plan was that once we'd gone through the material it would all be boxed up and shipped out to the USA to be held in the archives, thousands of miles from Scotland maybe never to be seen again. Back at home that evening, my wife said to me, "*You're very quiet. What's the matter?*" I was thinking we were about to lose everything. The factory was closing and the record of what was done there was about to go too. In the words of the Joni Mitchell song, "*You don't know what you've got 'til it's gone*", or in this case nearly gone!

My association with the South Queensferry plant in Scotland goes back to 5th October 1970 when I joined HP as an electronic design engineer to work on the Microwave Link Analyzer (MLA). I'd just graduated from Southampton University and earlier in the year some HP engineers came round the university looking for new recruits. In March, I came for an interview at South Queensferry. It was a long way from home, but that didn't seem to matter.

Hewlett-Packard was the world's premier test equipment manufacturer, and many viewed it as the Rolls Royce of electronics companies. I first became aware of HP at college in the late 1960s when a friend returned after the summer holidays with photocopies of the circuit diagrams from an HP oscilloscope (probably a 175 or a 140 series). We spent hours poring over the valve and transistor circuits trying to work out what the designers in Colorado Springs had done. Later at Southampton, I used a new 3420A Differential Voltmeter and was impressed not only by its accuracy but by the quality feel of the controls and the meter. Designing equipment like that was an opportunity not to be missed, so when I was offered a job I didn't think twice about moving north.

Today HP is recognised as a major supplier of computers, IT equipment and software, but in the old days it was famed for its technology and measurement instruments, the finest in the world particularly for RF and microwave measurements, signal analysis and frequency and time standards. That was HP's reputation in the golden age from the 1950s to the 1980s, and that's what drew me to the place.

On my first day at the factory, I was given the circuit diagrams for the MLA and told to familiarise myself. I remember being totally in awe of the analogue transistor circuit design and the inventive way it had been done. I wondered if I would ever be able to design anything like that. The MLA was only one of many remarkable products invented at South Queensferry over the years, innovation that was remarkable even by the standards of HP. There was some of the earliest use of digital electronics to generate and analyse noise plus the use of digital filters in the Dynamic Signal Analysis products of the late 1960s. There was the first direct digital waveform-synthesised signal generator in the 3770A Group Delay Test Set in the early 1970s, and then the first instrument in HP and probably the world to use an Intel microprocessor (3745 Selective Level Measuring Set). Alongside this, the tradition of outstanding circuit design in the MLA continued in several products for radio and analogue communications systems, some using in-house thin-film circuit technology.

In the emerging market of digital communications, South Queensferry engineers designed the world's first high-speed Pattern Generator and Error Detector (3760/61) and a few years later, their PCM Test Set (3779) made a breakthrough use of a menu-driven display and took digital signal analysis to new heights by using a proprietary HP computer chip, making it the most software-intensive instrument of its day. They stretched the technology of the time to its limits, and complemented this with some highly inventive ideas for solving measurement and communication problems. Not without reason, the Division was nicknamed inside HP, "The University of South Queensferry".

How did these innovations, and many more described in the following chapters, come about? Partly it was the ethos of the Company instigated by Bill Hewlett and Dave Packard. One of the design engineers from that time, David Dack, recalled the role of Peter Carmichael and Finlay Mackenzie, the inventors of the MLA. "*Finlay and Peter did exactly what Bill and Dave always advocated. First, collect a bunch of bright newly-qualified engineers – so new that they did not realise what they wanted to do was 'impossible'. Then let them get on with it, with minimal interference.*" Some from the R&D lab will remember the lengths we went to each year, trawling round the UK universities looking for the brightest engineering and computer science graduates. Of course, we just took it for granted at the time, but with hindsight it was undoubtedly one of the powerful drivers of the business.

Returning now to that archive meeting in the spring of 2010, and the question of what I could do to save something from this dying business. I took early retirement at the end of 2001 along with a hundred other old-timers, as the first wave of redundancies took place. At that time we employed about 2300 at the site. By 2010, following numerous further lay-offs there were just 200 staff remaining. As the place contracted, large areas of the factory were taken out of use. It was obvious that much in the way of documents and equipment had already been lost.

For the next couple of months I visited the factory once or twice a week and went through all the remaining documentation, and with the help of Mayumi Hepburn, copied what I thought might be useful as a record. There were thousands of photographs, mostly negatives, of products, people, events and views inside the factory. Fortunately, for nearly

40 years we had resident professional photographers at the site, so there were plenty of really good pictures. I scanned hundreds of these to provide a permanent photographic record. Coincidentally, another employee, Chris Burden, had found a large number of scanned photos (probably intended for *Readout* magazine) on two CDs, so between these two archives we have an excellent pictorial record.

Former colleagues brought round one or two old instruments and this formed the nucleus of a small collection. With some other instruments we found in the factory basement, we had a collection of around 14 items covering the 40 year history of the place. We put on a final product show in June 2010 and, with the help of Yvonne Mackie (site general manager), invited a curator from the National Museum of Scotland who agreed to take all of these instruments into the national collection. These are listed in Appendix 4.

In the final days, I found three ring binders abandoned in a cupboard that the product support department had compiled over the years. These had the product support plans and obsoletion notices for many of the instruments we designed. Information on dates, numbers produced and so forth. A goldmine that helped me produce the catalogue in Appendix 1.

My idea of writing up the history presented a challenge. How would I be able to combine a more general history of life and work at the factory, the social and business aspects, with the technical detail needed to do justice to the products? A former colleague, Magnus Hunter, suggested writing two books and that is what I have done. This book, Volume 2, describes the product families in some engineering detail, while Volume 1 is a more general chronological history about the business and social life at the factory, with much less technical information.

I should mention what's in this book and what's not. The product families are all those of the main Division at the site, which started in 1965 and went through until nearly the end. Originally it was the South Queensferry Division (SQD) and was Division 14 in HP's corporate organisation. After the operation focussed exclusively on Telecom Test, it became the Queensferry Telecom Division (QTD) and had the designation Product Line 63 (PL63). In the 1990s, it became Queensferry Telecom Operation (QTO) for a while, and then finally Telecom Network Test Division (TNTD), during which time it became part of Agilent Technologies in 1999 when the old HP was split up. The two other divisions on site, Queensferry Microwave Division (QMD) from 1985 onwards and the Telecom Systems Division (TSD) from 1994 onwards, are not included in this volume. The stories of these businesses and their products are covered in Volume 1.

The chapters about the product families are mostly self-contained, however some are intended to be read as a group, to avoid repeating a lot of background information. Chapters 2, 3 and 4 on Microwave Radio are one example, and the chapters on Telephone Line Analyzers and the RATES system are similarly related. Chapters 9 and 10 on Digital Transmission are intended to be read in sequence.

Have I viewed the past through rose-tinted glasses? I think I've given a fair assessment of the strengths and weaknesses of our products. Although well designed and built, they weren't all winners in the marketplace. However, the other facet of the development process was the disagreements and arguments between the various Division players in R&D and marketing. You might call it "creative conflict". Certain cliques formed around particular product areas and quite a bit of politicking went on promoting them, which could be a disadvantage if you were on the other side of the fence. On the other hand, one could say it showed that people cared and wanted to defend their ideas. As one former colleague put it, "*They simply cared passionately about the place and the work they were doing – apart from the obvious buffoons, that is!*" It was pretty competitive when it came to allocating resources, and projects sometimes got cancelled.

In the earlier years, many of these arguments revolved round the analogue versus digital communications market. In later years, another popular topic was the relative merits of stand-alone instruments versus network monitoring systems, which we called the "systems versus boxes" debate. There were also the inevitable personality conflicts between engineers with strongly held ideas and egos to match. Speaking to former colleagues while writing this book, it is surprising how quickly these old rivalries have resurfaced, even after 30 years or more. I've stayed away from this subjective stuff in the narrative as I don't want to reopen old wounds or offend people. In that sense, I've taken a charitable view of history.

So, what is special about the HP factory at South Queensferry in the context of Scottish industrial history? It was part of that enormous growth in high-tech industries in Scotland from the 1960s onwards – part of Prime Minister Harold Wilson's vision of a New Britain, "*forged in the white-heat of the technological revolution*"[1]. Many inward investment companies in the electronic sector set up in the central belt, attracted by government subsidies, and it became known as "Silicon Glen".

What set the HP factory apart from many of the other operations in Silicon Glen was the high level of product development – most of the goods produced at the site were also designed there and around 90% were exported. In later years nearly a third of the employees were involved with new product development. Scottish Enterprise saw it as a "Jewel in the Crown" of inward investment and I remember as PR manager in the 1990s, regularly hosting visits by foreign journalists and industrialists eager to see what we did. Finlay Mackenzie used to say that our product development capability gave us deep roots, meaning that it would be hard to uproot the operation and transplant it somewhere else. But even deep roots need nourishment and in the end a change in the market and lack of corporate interest meant the South Queensferry plant withered away and now it has gone.

Years pass, and for many who worked there it all now seems like another life as memories of it fade into history like a dream. History is made every day, but I feel sometimes there is too much emphasis on things that happened a very long time ago (how many more books on Mary Queen of Scots do we need?). The recent past gets less attention, particularly industrial history. I hope that HP South Queensferry will not be forgotten in this way. David Dack, one of the design engineers, summed it up with this quote from Shakespeare's Henry V:

> *Old men forget,*
> *Yet all shall be forgot,*
> *But we'll remember with advantages,*
> *What feats we did that day.*

---

[1] Speech at the Labour Party Conference, October 1963

As I write this Preface in the spring of 2013, Central Demolition is tearing down the buildings where these "Engineering Feats" were done. Soon there will be nothing left to mark the spot where so many original products were invented and manufactured. All that remains are living memories, the many photographs of life and work at the factory, this written record, and last but not least the products themselves as they are the embodiment of what the place was about and why we went to work there every day. In that respect we are fortunate as the instruments are mostly compact and highly concentrated examples of design and manufacture, and are a testament in themselves.

Of course, many of our earlier instruments are now completely obsolete as the changes in electronics and communications have been so dramatic in the last 30 years. Indeed, it would probably be hard to find telecom equipment anywhere in the world that they could still test! Instruments from the early 1990s onwards are compatible with modern networks and many are still in use today. All of these instruments had a part to play in the development of the worldwide communications network we take for granted. They were the tools of the trade and, although rather obscure and specialised, everyone has benefited from them indirectly.

So, this volume in the history of HP at South Queensferry is written for the engineer and I hope does justice to all these inventions. It is dedicated to the staff who designed, manufactured and marketed these products in South Queensferry for the global market. Their ingenuity and dedication made a great deal of money for HP and later Agilent Technologies. They were wealth creators, and it is something of a mystery why Agilent couldn't exploit this highly-tuned design and manufacturing capability, and just let it go. The site grew large and successful mainly through the ideas of these people. One can't help feeling that they, and their legacy, really deserved better.

I would like to thank the numerous former employees whose memories, insights and anecdotes have brought this volume to life. I have acknowledged the particular individuals at the end of each chapter. I would especially like to thank Finlay Mackenzie for his help and encouragement and for dredging up all kinds of memories from the "*recesses of his mind*". Finally, many thanks to Agilent Technologies for giving me a lot of access to information and equipment at South Queensferry in the final months, and allowing me to wander around the place at will.

Hugh Walker
Dunfermline
March 2013

Commitment to Innovation

*"We wanted to direct our efforts toward making
important technical contributions to the advancement of
science, industry and human welfare.
Right from the beginning, Bill and I knew
we didn't want to be a 'me-too' company,
merely copying products already on the market.*

*The key to HP's prospective involvement in any
field of interest is contribution.
To meet this objective, it is important that we put
maximum effort into our product development programs.
This means we must continually seek new ideas
for new and better kinds of products.
A constant flow of good new products is
the lifeblood of Hewlett-Packard,
and essential to our growth."*

Dave Packard

1

Chapter One

Introduction to Telecom Networks and Markets

Almost throughout its entire history, the South Queensferry site in Scotland was closely associated with telecommunications networks and measurements. It became the centre of expertise within the HP global corporation. The experience and range of knowledge in telecommunications technology would probably have rivalled that of the key manufacturers and research labs in the industry.

Over the years, tremendous expertise built up in a diverse range of topics: microwave radio (both analogue and digital), analogue and digital transmission, digital telephony, network management systems, data networks, network signalling, and high-speed fibre-optic transmission. Later, the neighbouring Queensferry Microwave Division became one of the key sites for the development and manufacture of test sets for the emerging mobile phone market, and the dominant supplier to cell-phone manufacturers.

None of this specialisation was pre-ordained when the factory was set up in 1965 – it was more an "accident of birth" that dictated its future.

In the earliest days of product development, firstly in Bedford (where HP opened its UK facility in 1962) and then in South Queensferry, there was a partnership with HP's Frequency and Time Division in Santa Clara, California. The first products to be developed were frequency counters (3734A and 3735A) and a frequency standard. HP UK was manufacturing frequency counters on behalf of the US division, and might well have pursued this line.

However, when the South Queensferry operation was established in 1965, two new development programmes were added following proposals by some enterprising new recruits from Edinburgh: the Microwave Link Analyzer (Chapter 2) and the Dynamic Signal Analysis products (Chapter 13). In the late 1960s, these became the main focus of product development at the site. As the story unfolded, the spectacular success of the

Microwave Link Analyzer convinced Bill Hewlett and Dave Packard that South Queensferry should focus its future strategy on the rapidly expanding telecommunications market, and so the destiny of HP South Queensferry was set.

The Evolving Telecom Network

The idea of this introductory chapter is to provide some background for the product family chapters that follow and to show some of the network models and market segmentation we used at the factory from the 1970s to the 1990s. The topic is a big one and comprehensive coverage is well beyond the space available here. Many books have been written about the architecture and evolution of telecom networks, and I have listed some references in the Bibliography in Appendix 2. These particularly cover the era and technologies which were important at HP South Queensferry in the latter 20th century.

It was an interesting period, starting in the 1960s and 70s when the network was almost entirely analogue, and then in the 1980s, migrating to digital telephony, switching and transmission, leading to the vast bandwidths of optical fibre which totally transformed the economics of telecommunications by the end of the century. By the late 1990s, which marked the final years for HP South Queensferry, another revolution took place when the Internet dominated the market and nearly all aspects of telecommunications migrated to packet-based transmission and away from the old circuit-switched philosophy.

Over all these years, and even to some extent in the 21st century, one can divide the network into four basic segments, and these remained broadly unchanged as the network migrated from analogue to digital.

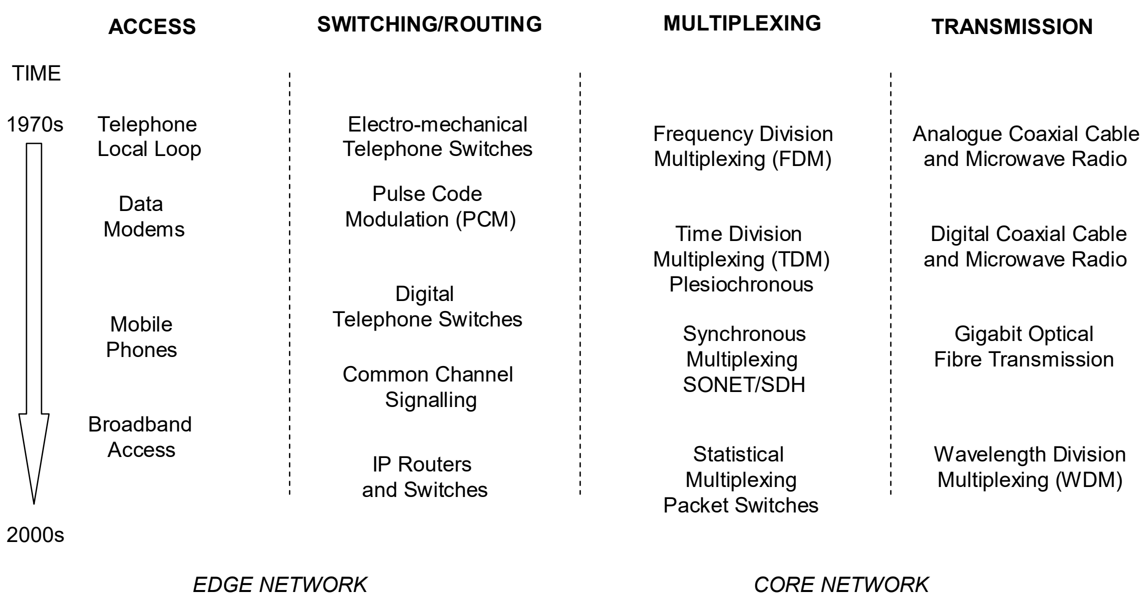

The **Access** segment connects the user or subscriber to the network. Traditionally, this was just the pair of copper telephone wires that connected the telephone to the exchange, known as the “local loop”. This is referred to as a 2-wire circuit and carries both directions of transmission on the same pair of wires. At the exchange, the directions are split and it becomes a 4-wire circuit with separate transmit and receive circuits. This is how the connection is handled through the main telecom network as shown here:

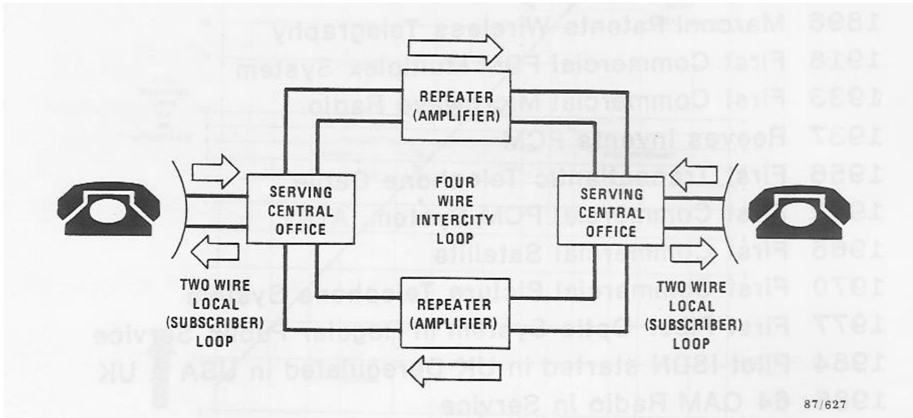

The same circuit would also be used for low-speed data using a voice-data modem, so that the data could be transferred over the switched analogue telephone network. Later, the same pair of copper wires was to be used to provide broadband access to the customer[1]. Download speeds up to 20 Mb/s were possible by using sophisticated digital modems with digital signal processing. The mobile phone network is another example of access, in this case the technology is wireless between the mobile and the base station over a distance of anything from 100 feet to a few miles. Once the signal reaches the base station it is handled thereafter by the fixed network.

The function of the **Switching/Routing** segment is to make connections between different users on the network. This may be tens of millions on a national network, or billions globally, so today the control of this section is very complex and relies on signalling networks and routing algorithms. In the 1970s, network traffic was almost entirely voice telephony and switching was electro-mechanical (the Strowger Exchange). This service was referred to as the Plain Old Telephone Service (POTS), and the system as the Public Switched Telephone Network (PSTN). By the 1980s, digital telephone switches were replacing the old mechanical space switches. This new generation operated with digitised telephone signals where the analogue voice signal was converted to digital form by Pulse Code Modulation (PCM). This topic of telephone switching is covered in more detail in Chapters 8 and 11, on PCM and Signalling.

Switching also has the effect of concentrating traffic. The domestic telephone is used fairly infrequently and spends almost all the time (maybe 95% or more) “on hook”. Business telephones may be used a lot more. Only when the user lifts the handset does the exchange respond and allocate switching to that telephone. Thus a local exchange, with say 10,000 incoming lines, may require only a few hundred outgoing lines to the network. This means that the more expensive long-distance transmission and switching equipment operates at higher traffic density and is more economical. The same concepts apply to both the early electro-mechanical and later digital telephone switches.

Similar principles also apply to modern-day Internet services based on IP (Internet Protocol) packets. The subscriber’s broadband access could have 10 Mb/s bandwidth or more, but normally it is only used intermittently and the traffic is a fraction of maximum throughput. The IP router combines and interleaves all these intermittent data streams from many customers and routes them as composite IP streams to the output ports. In this case

[1] ADSL – Asymetrical Digital Subscriber Line, asymetrical since the download speed is far higher than the upload speed, usually 10x faster. Subsequently, fibre-based VDSL with to 100 Mb/s download.

the traffic is combined statistically according to the demand. The customer's data might be just a few packets for an email message, or may be a continuous stream of packets for on-demand video. Each IP packet, whether it's an isolated transaction or a whole stream of data, has a specific address at the beginning of the packet (the header) which defines its final destination, and the routers in the network read this and automatically send it to the right place. The packets could take different routes through the network to get to the final destination, so there is no longer a direct mapping from the call connection to the physical circuit as was the case in the old telephone system. In IP, the connection only exists as part of the IP address and is described as a virtual connection, to differentiate it from the old physical connection of a telephone channel, analogue or digital.

The final segments in the hierarchical model shown above are **Multiplexing** and **Transmission**. Multiplexing simply takes the traffic concentration a stage further, in that it combines several lower-speed tributary data streams into one higher-speed stream, ready for transmission. Transmission systems generally give lower cost per channel, the more channels that are carried on a single cable or radio carrier.

From the 1950s onwards, there were enormous advances in the bandwidth of transmission systems and therefore the number of separate telephone channels that could be carried. This was driven particularly by the use of transistor electronics from the early 1960s onwards. Up until the 1980s, transmission systems used either microwave radio (typically handling up to 1800 channels per carrier) or coaxial cable (mostly up to 2700 or 3600 channels, though some systems went to over 11,000 channels per cable). Microwave systems used line-of-sight antennas with stations placed about 30 miles (50 km) apart. There's further information on this in Chapter 2. Coax systems required repeater amplifiers every few miles to boost the signal and compensate for the high frequency roll-off (attenuation) in the cable.

After the mid-1970s, these systems gradually moved over from the old analogue transmission to digital transmission. A transmission path can be thought of as having two dimensions – bandwidth (frequency spectrum) and time. Analogue systems use Frequency Division Multiplexing (FDM) and digital systems use Time Division Multiplexing (TDM), as shown here:

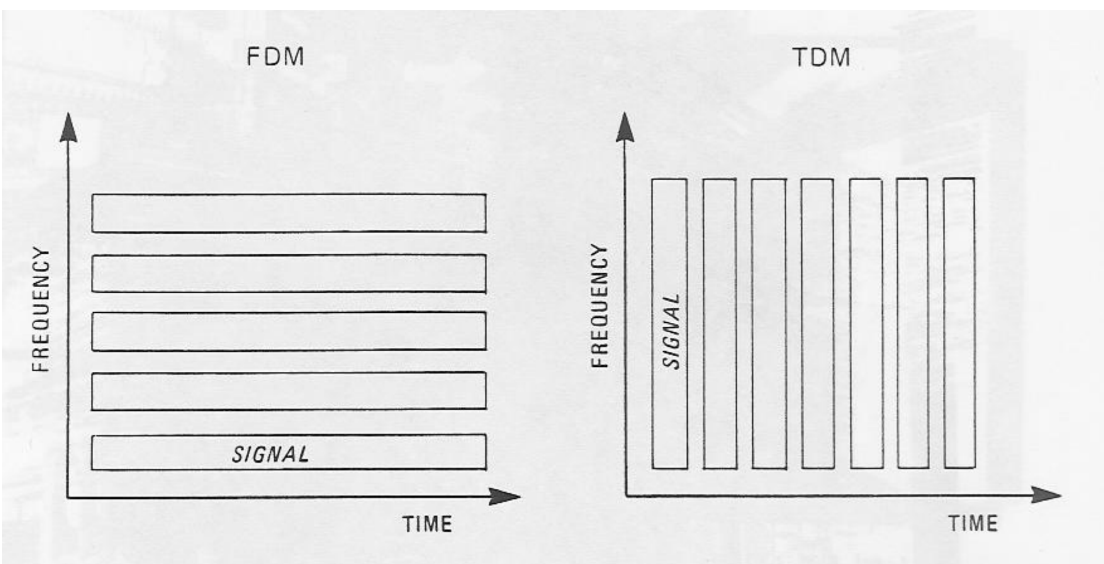

In FDM, the channels are allocated a portion of the frequency spectrum and have unlimited use of that band in terms of time. However, the transmitted signal spectrum must never lie outside this allocated band. In TDM, the individual channels are allocated the entire frequency bandwidth, but only for a limited portion of the time, i.e. a "time slot".

Here is another way to view the sections of the new digital telecom network. The triangles represent the quantity of equipment and investment in the network – greatest in the local-loop and exchange (or central office) and least in very high capacity fibre optic transmission. This representational diagram was popular at South Queensferry when discussing strategy, and had the nickname "*The Kissing Triangles*".

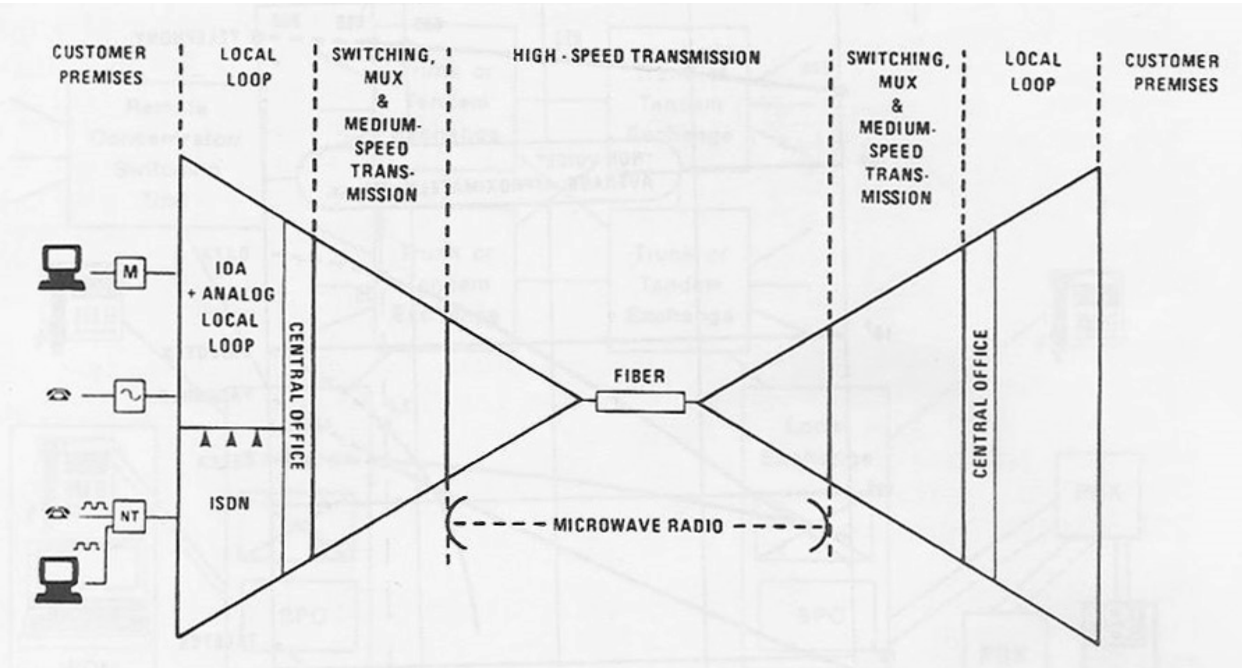

One of the discussion topics was about which segments it made sense for us to address. At first sight, the massive scale of the local loop and customer premises segment meant that a large amount of test equipment would be required. However, this requirement was mainly for simple and robust test sets making basic measurements, the kind of thing a telephone engineer took in his van[2]. On the whole, HP as a company specialised in more sophisticated equipment that commanded higher prices, equipment that was more difficult to make so few could compete. Our resulting higher cost structures made the low-end market unattractive.

Although, in later years, South Queensferry did start to compete in these areas, partly through acquisitions, the home designed products were almost exclusively focussed on the medium and high-speed transmission segments of the "Kissing Triangles". This started in 1967 with the first Microwave Link Analyzer, and carried on right through to the final Digital Transmission Analyzers operating at Gigabit rates, 35 years later.

[2] In the 21st century, the move to digital subscriber services like broadband access and Ethernet, has meant that the kind of measurements once done at the core of the network have now moved to the edge, so the "telephone engineers" test set is far more sophisticated but still has to be robust and portable.

In-Service and Out-of-Service Testing

South Queensferry's first telecom test set, the Microwave Link Analyzer, was used to align microwave links for best performance. It did this by applying specific test signals and measuring any resulting link impairments at the far end (as described in Chapter 2). During equipment manufacture and installation this presented no problem, but once in revenue-earning service, the link had first to be taken out of service and the telecom traffic switched to a standby or protection channel or rerouted. There was no way of accurately adjusting the link with the traffic present. This is called Out-of-Service Testing.

The same principle applied to the test of digital transmission systems up until the late 1980s. In this case the traffic signal was removed, a pseudo-random test signal applied to the system to simulate traffic and the received signal checked bit-by-bit for errors. This was the normal method for measuring the error performance of first generation digital systems conforming to the Plesiochronous Digital Hierarchy (PDH) standard, as described in Chapter 9.

The only Queensferry product from that era that did in-service testing (that is with revenue-earning traffic still present), was the 3745 Selective Level Measuring Set (SLMS) of 1975, and its successors. It was designed specifically for measuring the live customer traffic and the various control signals in a working analogue FDM transmission system. As described in Chapter 5, its early use of microprocessor control gave it the power to scan large numbers of signals on the system at speed, so a good assessment of system performance was possible in-service.

In-service testing or performance monitoring was always the preferred method for network operators or service providers, since the measurements didn't interrupt the customer traffic, and degraded performance could often be detected early, before it caused disruption. Increasingly, operational equipment had built-in performance monitoring and alarms that did this job, but there was also a role for external network monitoring systems. These checked performance and traffic characteristics through specific monitoring points on the operational equipment, collected data over time and presented it via computer analysis. South Queensferry was active in this market and the first monitoring systems were built in the 1970s using the 3745 SLMS mentioned above. Later, a highly successful monitoring system for No.7 signalling was introduced in the 1990s (see Chapter 11) which took the idea further by analysing traffic statistics and specific messages transferring through the system.

The Fibre Optic Era Arrives

By the late 1980s, the global telecom network was moving rapidly to completely digital transmission and switching. The first generation digital transmission systems using coaxial cable and microwave radio were being replaced by high-bandwidth fibre optics. Single-mode, long-wavelength fibre cables, high-speed laser transmitters and new high-technology integrated electronics, transformed the telecom business. The economics of long distance transmission changed radically with much higher bit-rates in the Gigabit per second range made possible with the low loss and wide bandwidth of fibre optic cables. HP introduced a new range of optical test gear (mostly developed in Germany and Santa Rosa, California) describing it as Lightwave Instrumentation.

Up until the late 1980s, the standard digital transmission system used PDH multiplexing which operated at 2, 8, 34 and 140 Mb/s. A different PDH system was used in North America, the primary rate being 1.5 Mb/s (DS1) and the main transmission rate being 45 Mb/s (DS3 being 28 DS1 tributaries). Multiple DS3 streams were combined for trunk transmission, a popular rate being 135 Mb/s (3 x DS3). Here's a Division slide from the mid-1980s – very familiar to the R&D and marketing staff at South Queensferry, as it was the "bread and butter" of our business. It shows the European and North American (Bell Standard) PDH rates. Note that the only traffic envisaged at this time on the network comes from audio telephony (and data modem signals carried over the telephone channel).

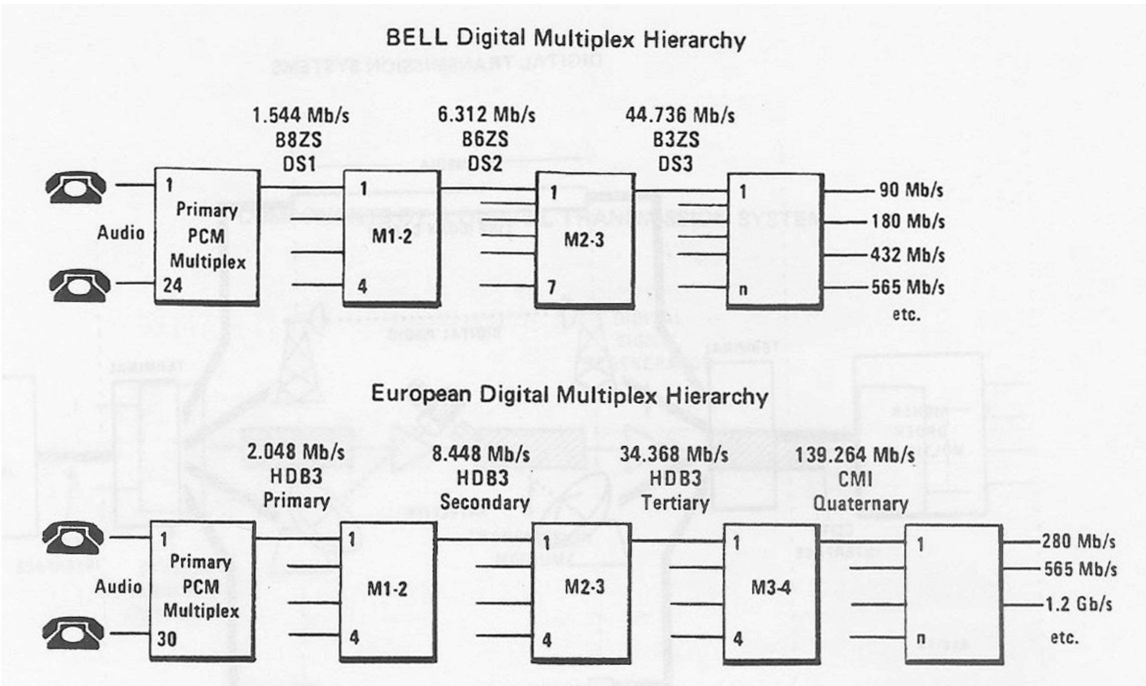

Complementing the optics development in the late 1980s, a new digital transmission standard began to be used in the industry. This second generation used synchronous multiplexing which allowed much easier access to lower-rate tributary streams, as explained in Chapter 10, so operators could manage bandwidth and traffic more efficiently. This new standard was called the Synchronous Digital Hierarchy (SDH), and in North America the Synchronous Optical network (SONET). The standards included both electrical and optical specifications as it was expected that fibre optic cable would be used.

Bit rates were much higher than PDH, with the main hierarchy rates being 155/622/2488 Mb/s and 10 Gb/s by the late 1990s. The other key difference was that the synchronous multiplexing frame structure had much more comprehensive error checking, alarms and other management/control information built in.

From 1990, Queensferry started to introduce a new family of Digital Transmission Analyzers capable of testing to this new standard with much higher bit rates and operation to the new frame structure with optical interfaces. Along with some new Gigabit pattern generators and error detectors for measurement of optical transmitters and receivers, by the late 1990s the Division had a remarkably successful portfolio of products focussed on this new era, as described in Chapter 10.

Network Signalling

Signalling is the means by which a call is set up through the network. In the old days of POTS, it was from one telephone to another, and relied on the dial pulses generated by the rotary dial on the phone. The whole exchange network was largely electro-mechanical, the dial pulses acting directly on the relays and rotary switches. The signalling was closely associated with the actual switching mechanism, and was known in the industry as Channel Associated Signalling (CAS). Quite a lot of variants of this principle were developed from the 1930s onwards, particularly in the post-war period when long-distance subscriber dialling became available, since the dialling information needed to be sent over transmission facilities rather than copper wires.

From the late 1960s, a new generation of telephone switches entered the network using computer control, known as Stored Programme Control (SPC). Initially these were relay switches controlled by computers, and later digital switches in the 1980s. Since the computer control was now functionally separated from the telephone switching, the computers communicated through dedicated data channels to send signalling commands between the exchanges. This was known as Common Channel Signalling (CCS) and is described in greater detail in Chapter 11.

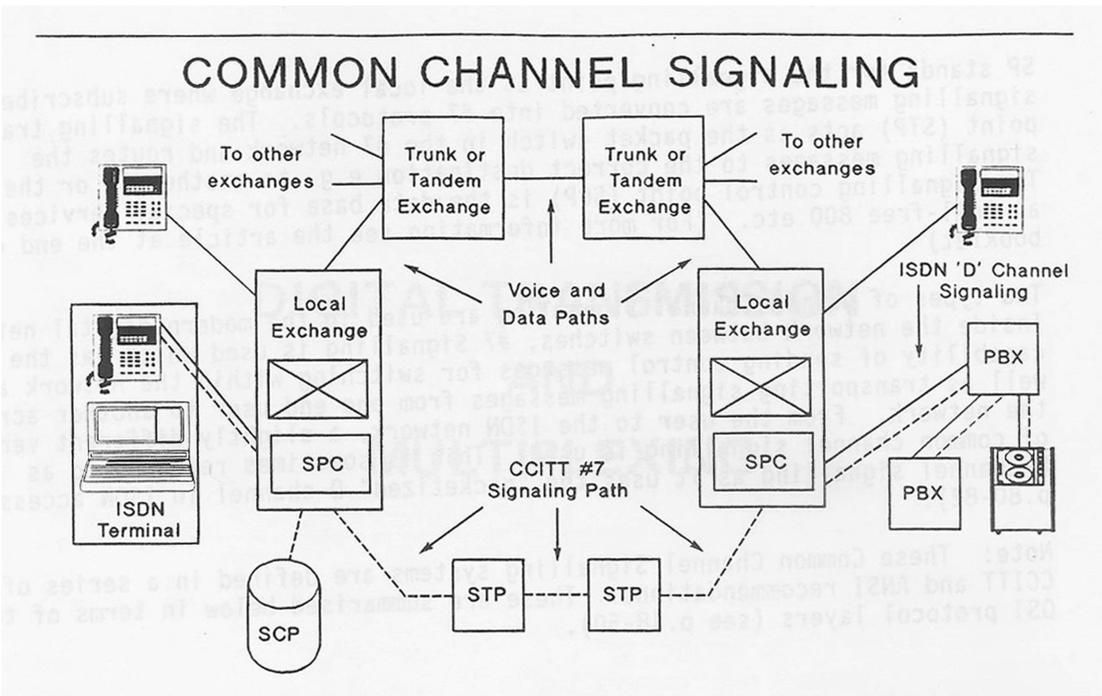

Whereas the earlier CAS systems used a multitude of different standards, the new CCS systems rapidly converged on a single global standard called Signalling System No.7 (SS7), from around 1980 onwards. The market was large and the possibilities extensive since the new system allowed all kinds of other operations apart from the conventional telephone call setup. Features like Freephone 800, premium rates, card billing, voice messaging and other advanced services were introduced, which would never have been possible with the old CAS approach.

South Queensferry management became interested in this emerging market and in the late 1980s introduced a family of signalling test sets which could decode these messages as well as generate signalling traffic. The market became even larger in the 1990s with the growth of mobile phone networks which also relied heavily on SS7 for network control. As described in Chapter 11, the Division eventually produced a winning product that had much success in the 1990s, and led to the development of SS7 in-service monitoring systems.

Telecommunication Standards

International standards have been a crucial part of telecommunications industry, ensuring that traffic could be handled seamlessly between different national networks and operators, as well as guaranteeing equipment from different manufacturers, for example switches and transmission systems, could inter-work correctly. Standards have also been established over the years for the expected level of network performance (e.g. noise, distortion, digital error rate and so on) particularly on international connections.

Test and measurement equipment is a key part of this since it will be used to check that equipment conforms to the inter-working standards and that overall network performance objectives are being met. South Queensferry therefore had a strong interest in Telecom Standards. In fact it was the bedrock of all the product development work. A significant amount of shelf space in the Division's library was devoted to standards documentation from many sources which gradually accumulated over the years. This was markedly different to most of the other instrument divisions in HP which developed products driven more by technology and competitive specifications. If you built oscilloscopes or spectrum analyzers, the most obvious strategy was to make a better one cheaper.

Of the various standards bodies, the primary interest for Queensferry was CCITT, The International Telegraph and Telephone Consultative Committee. There was also a sister organisation, CCIR, which focussed on radio systems. Both were part of the International Telecommunications Union (ITU) based in Geneva, which represented the world's telecommunications authorities, much as the United Nations does in politics.

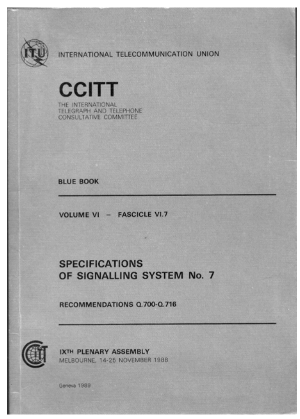

The CCITT Recommendations were issued every four years in dozens of books identified by the colour of the cover. Thus there were Green Books, Red Books, Blue Books and so on. It was a familiar sight to see these colourful books on the desks of engineers in marketing and R&D. The last complete set of printed standards was the Blue Books in 1988, an example of one is shown here. Thereafter, the standards were issued and updated as required, and today are available on-line.

These CCITT Recommendations, later renamed ITU-T (ITU – Telecommunication Standards), were sometimes loosely called the "European" standards, although they were used by most countries in the world, as intended. The exception was North America where a very different set-up prevailed. This went back to the days prior to divestiture in 1984, when the whole network operated as a monopoly under AT&T and Bell Telephone. It was a vast network and it was quite practical to specify its own standards independent of the rest of the world. In contrast to the CCITT "European" standards, these were referred to as the "Bell" standards[3]. They were so different, that network equipment and test sets needed to be designed specifically for the North American market.

Some advocates of the CCITT standards looked down on the US Bell System as the network did have quite a lot of odd anomalies that accumulated over the years due to the way it evolved. CCITT systems were probably cleaner because they had been thought through more, although it wasn't politic to mention that! By the 1990s, the difference began to disappear, and many of the old Bell standards became part of the CCITT documents. No.7 Signalling became universal, and the new synchronous multiplex and transmission standards, SDH and SONET mentioned earlier, were intentionally unified to a large degree.

The importance of these standards can't be overstated. Unless test equipment conformed to the requirements, it would be difficult to sell in the market. South Queensferry was deeply immersed in this area, and several engineers from marketing and R&D got involved in the CCITT Study Groups which developed and revised these standards. In particular, the Division contributed to the CCITT O-Series Recommendations which defined test equipment and test methods.

Market Strategy

Writing about telecom market evolution from 1970 to 2000 is easy to do retrospectively. The various technical developments and market changes now look fairly logical and the pieces all fit together like a jigsaw. But imagine yourself going back to the early 1980s and trying to decide what would be the winning technologies and how fast the market would change. You needed a "crystal ball" combination of judgement, knowledge and quite a bit of luck to back the right future products.

It was certainly the case that the more knowledge you had about the market, the more likely your judgement would be on the right track. So, the Queensferry engineers in R&D and marketing spent a good bit of time talking to and listening to customers all over the world to get a feel for where things were going. The aim was to identify some key industry experts and build a relationship with them. In management parlance it was called "*analysing the voice of the customer*". Finlay Mackenzie, marketing manager at the time, recalled some advice from Bill Hewlett about future product strategy: "*Listen to customers not the sales folks because they are only interested in what they can sell today, whereas customers will tell you about future needs.*"

Other sources of "knowledge" came from activities in standards committees which the Division participated in. If you were already a key player in a market, particularly a new and emerging market, you probably had quite a lot of information on the size and growth of the market and what competitors were doing. Large supply contracts won and lost was also a good barometer of the business.

Here are some of the questions we tried to answer. What is the size and growth of a particular market segment? Which products will give the best return on investment?

[3] Key documents were the Bell System Technical References (BSTRs) which defined operating and interface specifications for network equipment as well as performance criteria and test methods.

Should we focus on the large analogue communications market and try to take market share from competitors, or should we go for the smaller and faster growing business in emerging digital communications? Should we develop products for the European or the North American market, or try to do both with a common platform?

Another key question was to do with customer market segments. We saw earlier the segmentation by network section – the "*Kissing Triangles*" model. The other view is how the customers themselves could be segmented, depending where they were in the economic model or "food chain". Where would the Division have most success?

Broadly, the customers could be segmented into three groups:

- The **End Users** who were the customers for telecom services. They could be private subscribers all the way to major commercial organisations.
- The **Service Providers** were the telephone companies who sold the services to End Users.
- The **Equipment Manufacturers** who designed and built the telecom equipment used in the network, and sold it to the Service Providers.

Here's a market model used at the Division in the late 1980s, showing the interaction between these groups and the money that was flowing on an annual basis.

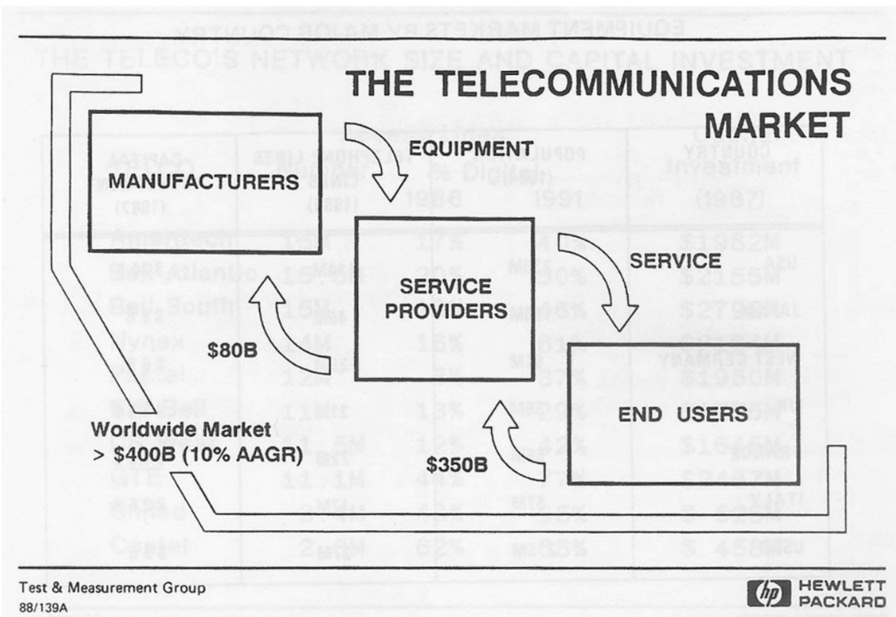

These are massive numbers and the test equipment was a tiny fraction of them, but there was potentially more business to be done with service providers and their customers, than manufacturers. On the other hand, HP's focus on high-end instrumentation played better with the equipment designers and manufacturers. However, by the 1980s service providers were investing in network monitoring and automatic test systems, so there was potential if the markets were accessible. One could describe this as a "top down" approach.

All these questions were hotly debated during many hours of meetings at the factory over the years. Much of the interest was in developing market models using both the top-down and bottom-up methods. The latter involved trying to build up a particular market segment from known pieces of data, for example we did $5M of business in this segment last year and believe we had 20% market share. Who are the other competitors and what business did they do?

I remember doing this exercise in the late 1970s in a meeting room overlooking the Forth Bridge as we tried to fill in a matrix of market segments versus HP and competitors. Finlay Mackenzie referred to the process as "*Rows and Columns*", or in his Hebridean accent, "*Rowsh and Coyl-yumsh*". Commenting on the strategy booklet that resulted, Finlay recalled, "*I think the page I derived most satisfaction from was the one that had a matrix of Customer Needs – Current and Future, for each segment.*"

This diagram, which we used around that time, shows the application segments sometimes used in these strategy discussions:

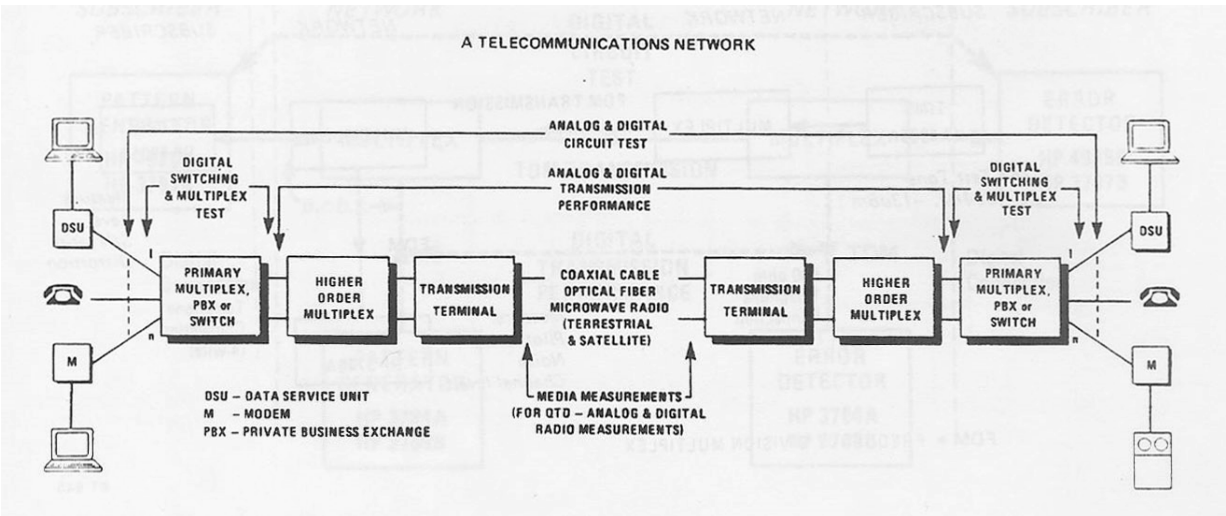

Another key strategy event in the 1970s and early 1980s was the annual off-site meeting at the King Malcolm Hotel just across the Forth Bridge in Dunfermline, usually lasting a couple of days. Here the product development plan for the next five years was thrashed out. There was inevitably quite a bit of axe-grinding as various players promoted their particular interests, sometimes using rather "selective" market information. In the end though, we had to have a priority list based on financial return and growth. If things were unresolved at the end of the second day, it was a question of whether to order sandwiches. "*No, no, not the sandwiches. Let's make a decision. I want to get home for my tea!*" I remember on one occasion when such an impasse had been reached, Finlay got up and wrote 1, 2, 3, 4, 5 etc. against each of the projects on the overhead slide in his order of priority. "*You can't do that*", everyone said, to which his response was, "*Well I just have – tell me what's wrong with it!*" Years later, although the meetings no longer took place at the hotel they were still referred to as the "*King Malcolm Meetings*" and everyone knew what it meant.

Memories of most of these meetings and discussions have long disappeared in the mists of time, but the result of these deliberations is evident in the products we finally developed. All these are described in the chapters that follow. Some were very successful, but others failed to live up to expectations. Product strategy was far from an exact science. Products could turn out to be too early or too late into a market, and sometimes we didn't fully understand customer requirements. They could be "blind-sided" by unexpected changes in the market and technology, or the unseen entry of competitive products. The risk of this was greater if the product was too long in development, which was something the Division could control. Much effort went into reducing time-to-market by improving the product development process as the years went by (see Volume 1, Chapter 7).

Summing up

For several years in the early 1990s, South Queensferry published a short-form catalogue featuring all its products along with communications products from other HP Test and Measurement divisions.

It had the tagline:

"*Helping you optimize network performance*".

It included a diagram (shown on the next page) that neatly summarises the telecom network discussed in this chapter and lists the measurements, many of which are described in more detail in the product chapters that follow.

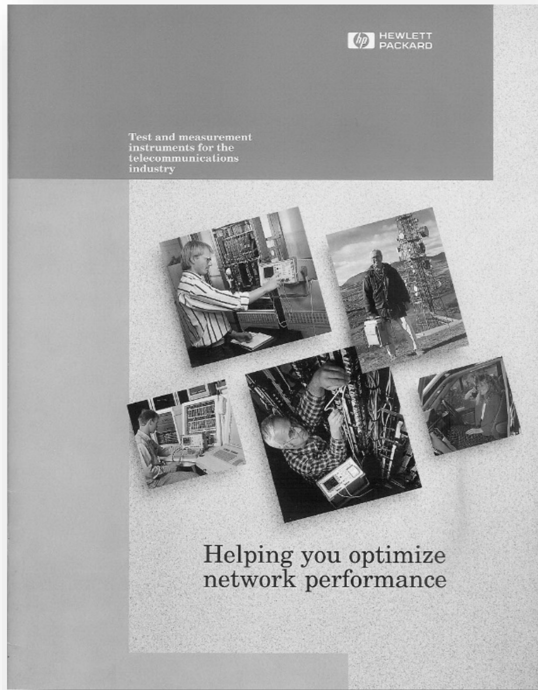

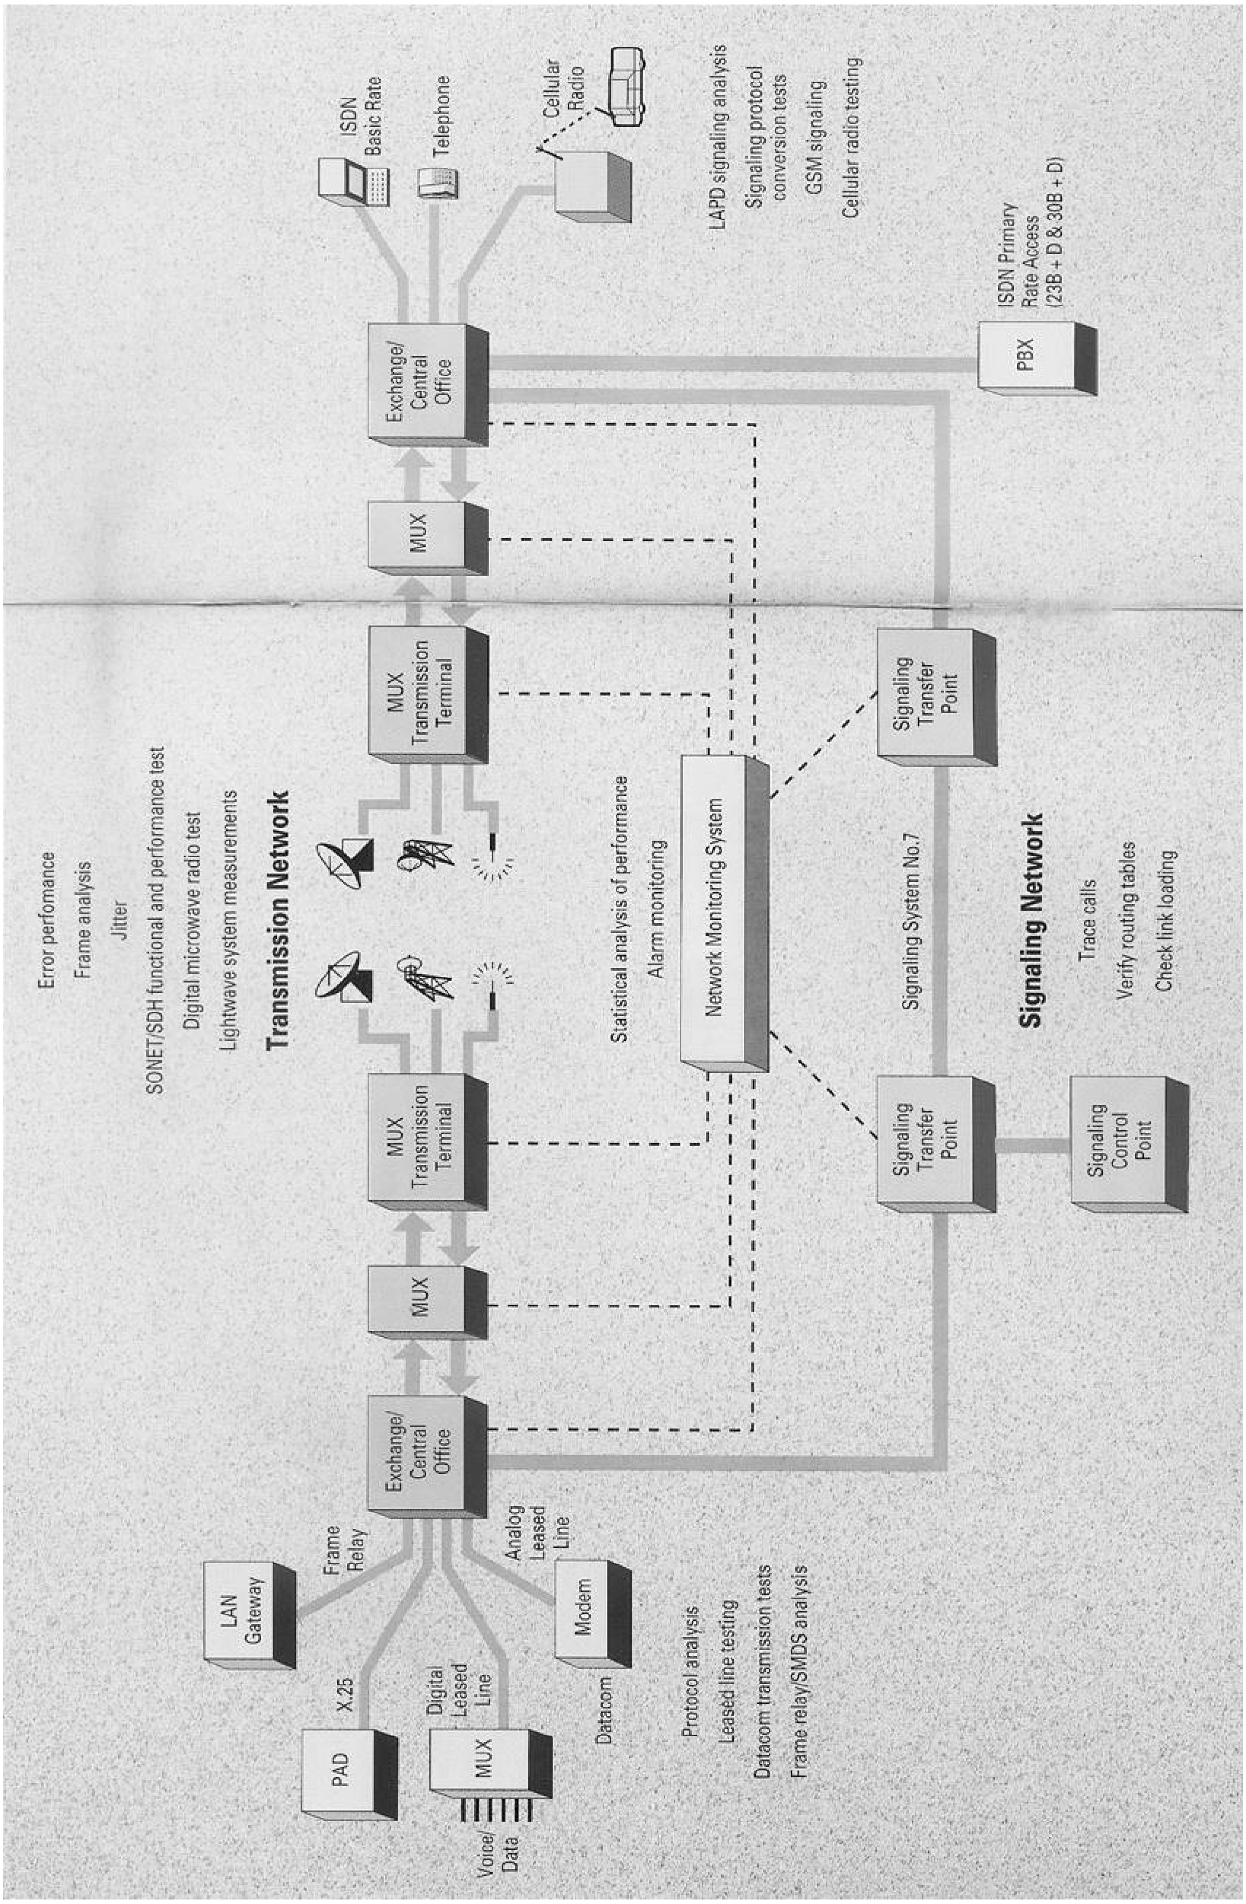
Error perfomance
Frame analysis
Jitter
SONET/SDH functional and performance test
Digital microwave radio test
Lightwave system measurements
Transmission Network
ISDN Basic Rate
Telephone
Cellular Radio
LAPD signaling analysis
Signaling protocol conversion tests
GSM signaling
Cellular radio testing
Exchange/ Central Office
MUX
MUX Transmission Terminal
MUX Transmission Terminal
MUX
Exchange/ Central Office
PBX
ISDN Primary Rate Access (23B + D & 30B + D)
Statistical analysis of performance
Alarm monitoring
Network Monitoring System
Signaling Transfer Point
Signaling Transfer Point
Signaling Control Point
Signaling System No.7
Signaling Network
Trace calls
Verify routing tables
Check link loading
LAN Gateway
Frame Relay
X.25
PAD
MUX
Digital Leased Line
Analog Leased Line
Modem
Datacom
Voice/ Data
Protocol analysis
Leased line testing
Datacom transmission tests
Frame relay/SMDS analysis

2

Chapter Two

The Microwave Link Analyzers

In the summer of 1958, Peter Carmichael and Finlay Mackenzie joined a new group at Ferranti Ltd. in Edinburgh (subsequently BAE Systems and now Selex at Crewe Toll) developing a ground radar system for the Bloodhound Missile. After completing the radar in the early 1960s new military contracts were hard to come by, so Ferranti decided to use the expertise to develop commercial microwave links. At the time, these were used for long-distance trunk transmission of multiple telephone calls, as many as 1800 telephone channels per microwave carrier. The same systems were also used for TV distribution round the country. One major microwave radio station was located at Kirk o' Shotts in Lanarkshire.

Peter experienced difficulties testing his designs and began to realise there was a lack of dedicated commercial test equipment to carry out measurements on microwave links, and particularly the modulators and demodulators he was working on. Furthermore, he realised that with the increasing number of links in use throughout the world, an efficient system that would allow engineers to readily assess the operation of their links and adjust them for optimum performance, had a large potential market. This he discussed with his colleague Finlay.

By 1964, Peter had developed his ideas sufficiently to put a proposal for a new microwave link tester to his employers, Ferranti. Perhaps because it was outside their mainstream business, they declined to take it forward. He needed a more receptive audience, and who better than the world's premier test equipment manufacturer, Hewlett-Packard. By strange coincidence, HP was at that time in the process of relocating from Bedford to South Queensferry. They placed an advert in the local paper and Peter went for an interview. Asked if he had some new ideas, he said "*Yes, a test set for microwave radio links*". Meanwhile, he arranged an interview for Finlay a week later at the Hawes Inn one Saturday morning. Peter joined HP in May 1965, the month work started on building the new factory, and Finlay a couple of months later.

So, what is a microwave radio link and what was it that needed the special test set Peter had invented?

Anatomy of the Microwave Link

When someone makes a long-distance telephone call to another town or city using an STD or area code, the call is directed onto to the trunk network rather than being handled entirely at the local exchange. In the early days, the trunk network consisted of hundreds of separate telephone wires inter-connecting the trunk exchanges. By the 1960s and 70s, this trunk network was increasingly implemented using frequency division multiplexing (FDM) systems whereby the individual telephone channels were allocated a specific frequency in the spectrum, a bit like old-fashioned radio broadcasts on the AM and FM frequency bands. Typically up to 1800 telephone channels were combined in this way, later rising to 2700 channels.

This composite multiplexed signal was then transmitted directly through a single coaxial cable buried in the ground, or was modulated onto a microwave carrier and transmitted across a microwave link. Both systems were used extensively in the analogue communications era prior to the 1980s and the advent of digital fibre optics. Cable meant you had to have access to dig a trench, and also the signal needed to be amplified every few miles. Microwave radio on the other hand simply required a radio mast and equipment every 30 miles or so, and avoided the need to access the land in between.

The microwave link usually operated in the 4 GHz or 6 GHz band (later also 11 GHz), and was a Line of Sight (LOS) system whereby the transmitter antenna (a parabolic reflector) was beamed directly at the receiver parabolic antenna about 30 miles away. It had to be high enough to avoid obstructions, so some radio stations were placed on top of hills, or structures like the Post Office Tower in London. In the UK, most trunk connections could be implemented with three or four 30 mile hops, whereas in North America where microwave links dominated, enormous distances had to be covered with many hops.

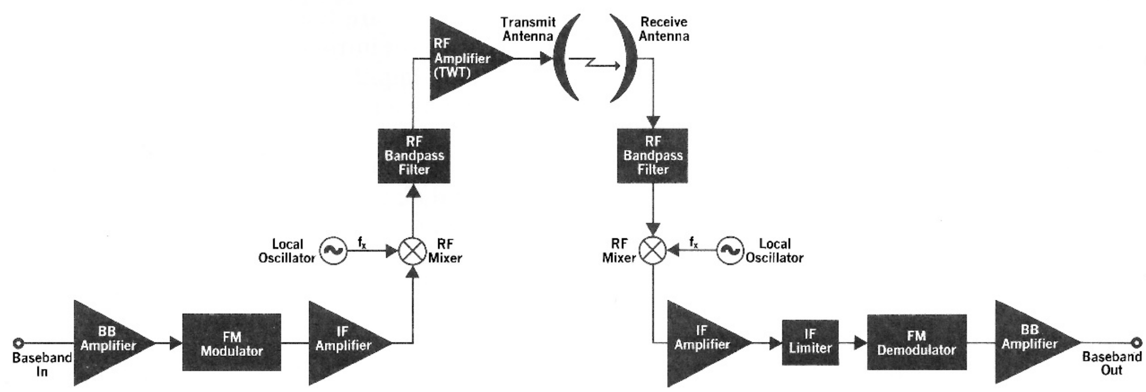

This diagram shows the main elements of a microwave link. The baseband input is the multi-channel telephone signal described above. It modulates the frequency of an oscillator at around 70 MHz (the intermediate or IF frequency) and the signal is then up-converted to the microwave transmitter frequency. At the receiver, the signal is down-converted and the baseband recovered using a frequency demodulator. As mentioned earlier, the largest distance between the transmitter and receiver is typically about 30 miles (50 km), so for long distances, intermediate stations called repeaters are required. These simply receive the signal, amplify it and retransmit it. They do not recover the baseband, so do not have the modulator and demodulator.

Because of the complex signal being transmitted, the various elements in the system need to have a very tight specification. The modulator and demodulator need to have an extremely accurate linear relationship between the baseband signal and the corresponding frequency deviation and vice versa. The IF and microwave sections all need to have a very flat frequency amplitude response as well as good phase linearity (flat Group Delay). Because this is an analogue system, all the individual impairments gradually accumulate, and over a long system could lead to an unacceptable degradation in quality. For telephony this would be experienced as a high level of background noise on a telephone call, while for TV signals the picture quality and colour rendition would be poor.

The Measurement Challenge

Anyone familiar with electronic measurements will realise that measuring and adjusting these individual system components in the microwave link will be difficult to do accurately with general purpose test equipment even today, let alone in the early 1960s. This was the frustration faced by Peter Carmichael designing links at Ferranti.

Furthermore, it is one thing to measure some circuitry on the test bench where signal generator and the measurement receiver are sitting together and can if necessary derive a local reference signal – the measurement receiver usually compares the received signal with the unimpaired test signal to evaluate the impairments in the item under test. But what happens when the signal generator is tens or hundreds of miles away from the measurement receiver at the other end of an installed link?

This last issue is a common problem with measurements on telecom systems, and many of the products developed at South Queensferry over the years had to find a solution to the "end-to-end" problem of remote transmitter and receiver. In effect, the receiver has to reconstruct or synchronise a reference signal from the incoming measurement[1].

A final challenge for the designers of the Microwave Link Analyzer was to ensure the measurement specification could be guaranteed even if a random pair of transmitter and receiver test sets operated together. It is quite easy to optimise the back-to-back performance of a specific pair of instruments on the bench, but out in the field it is more than likely two units will have to inter-work that have never been tested together before. Calibration methods were developed at South Queensferry that allowed HP to claim the published specification would be maintained whatever transmit test set was used with whatever receiver. That was a valuable reassurance for microwave link engineers.

[1] In analogue test sets this was typically done by locking on to the incoming signal with a narrow-band phase-locked loop which generated a stable reference signal for comparisons. In a few cases, a fixed reference was transmitted as part of the measurement sequence which the receiver used as the reference. In digital test sets, usually the transmit pattern generator circuit was replicated in the receiver as a reference and then synchronised to the incoming data stream for bit-by-bit comparison. The test set would also have to recover its own timing clock signal from the incoming data.

The HP Solution

In any general measurement situation, it is a good rule of thumb that the system being used to make a measurement should be ten times better than the item being tested. This ensures that the result is an accurate measurement of the item under test with a minimal contribution from inaccuracies caused by the instrument.

There are two ways to achieve this. The simplest solution conceptually, is to use a set-up similar to the item being tested, but ensure it is very well calibrated as a reference standard. This is the approach often taken in standards labs, and is sometimes referred to as the "gold standard". The other approach is to take a fresh look at the measurement and the source of errors and try to devise a measurement method that inherently eliminates the measurement error. In this case the instrument will probably not be much like the item under test. HP had a long history right from the earliest days in 1939 of designing test instruments that did just that. By a clever combination of mathematics, system design and circuitry, they produced a whole family of products that were reliably much more accurate than other solutions on the market.

This was the approach Peter took for the new Microwave Link Analyzer (MLA). In October 1965, he set off for the USA with a proposal for the new instrument. He visited many of the key manufacturers, design labs and operating companies, accompanied by HP's sales director, Bob Brunner. After the trip, Bob commented that it was the most promising market survey he had conducted for 15 years. A star was born. While in the USA, Peter was introduced to Barney Oliver, Vice President of R&D who ran HP Labs. Barney had "*a few minutes to spare*", but was so impressed he spent two hours with Peter.

This trip helped Peter and his colleague, Finlay Mackenzie, define the features and specification of the new instrument and soon work began on the design. At that time, the new HP factory was still under construction, so the R&D team were offered space in the council chambers in South Queensferry. During this period, more customer visits in Europe and USA resulted in further refinements to the overall product specification. Finlay recalled a pivotal visit he and Peter made to Telettra in Italy, at the time one of the leading microwave link manufacturers. They asked for many more features, and two of the Telettra research team visited South Queensferry to make sure their ideas were incorporated! Finlay considered their inputs key to the success of the MLA. Later, on the product introduction tour in the USA, Finlay remembered the positive reaction from the engineers who couldn't believe that their needs were so well understood:

"*I just told them we had listened carefully to many engineers before designing the block diagram. Telettra had made us incorporate extra features. The wide 50 MHz sweep they demanded was a challenge and Peter and I wrestled with the complexity of the change but we decided we had to do it.*"

A complete technical description of the MLA is well beyond the scope of this account, but to give a flavour of its innovative design, here is an overview of a key part of the system.

As mentioned earlier, the new MLA needed to achieve a specification far better than the microwave links it would test in order to provide reliable measurements. A typical microwave link had an operating bandwidth of around 30 MHz to transmit the 1800 channels of telephone traffic, so the MLA needed to measure and display an accurate representation of the impairments over that bandwidth. The principal used was to sweep the carrier frequency across the active bandwidth (with a low-frequency 70 Hz sine-wave) and superimpose a much smaller high-frequency signal called a test tone. The idea was that any impairments in the microwave link or its individual components would cause variations in the amplitude and phase of the test tone as the carrier was swept across the band. The receiver would then demodulate and isolate the test tone and measure the phase and amplitude variations to create a swept display of the impairments.

For this idea to work, the instrument design had to ensure minimum interaction between the low frequency sweep and the test tone so that, back-to-back, the MLA would give a flat response. The inventors came up with a beautifully elegant solution shown here:

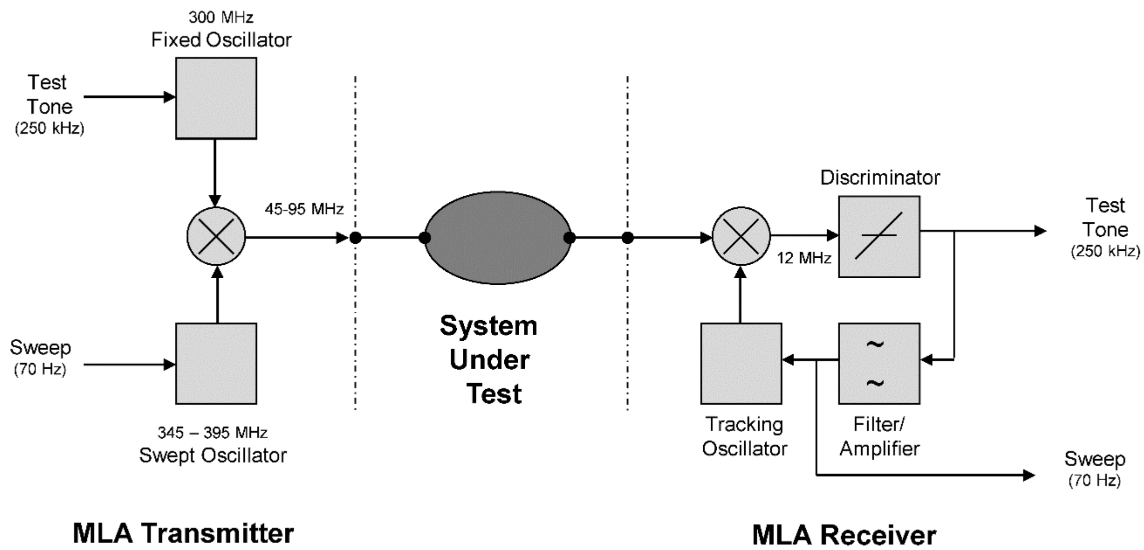

By frequency-modulating two independent high frequency oscillators with the sweep and test tone, and then mixing the outputs, the interaction between the two test signals is minimal. The MLA receiver uses another clever idea of a tracking feedback loop which forces the tracking oscillator to exactly follow the sweeping carrier signal. This means the frequency demodulator or discriminator only sees the test tone so again the interaction between test tone and sweep is minimised. Because the test tone is more than a thousand times higher in frequency than the sweep it is easily filtered out and doesn't modulate the tracking oscillator. The recovered sweep signal is used to scan the receiver display and the recovered test tone is processed to provide the measurements. Notice that the only connection necessary between the two instruments is through the measurement path itself, so the end-to-end problem is solved.

This simple but effective system design was at the heart of the MLA. So good was the inherent performance that the same basic design was used on later versions of the instrument and, with further enhancements, it soon established itself as the global industry standard for microwave link measurements and was never really surpassed. The design team incorporated a whole range of additional measurements into the instrument, often by clever reuse of circuits, so that the new test set combined most of the measurements the engineer needed to test a link, all in one pair of instruments. At the time, that was a remarkable achievement.

A World Beater – The 3701A MLA System (NMS T.2010.65.1-3)

In a memo to the North American field sales engineers dated 20 September 1967, Jim Cochrane launched the new instrument and recalled the tour made by Peter Carmichael and Bob Brunner a couple of years earlier:

"*Visits at Farinon and Lenkhurt in San Carlos, along with visits at Collins in Dallas, General Electric in Lynchburg and Bell Labs at Andover, were the basis for a black box with all the 'right knobs'. Continuing correspondence with these companies kept the R&D efforts at HP Ltd. in step with the requirements of Microwave Link Analyzers in this country …. and now HP Ltd. is ready to demonstrate their Model 3701/2/3A Microwave Link Analyzer.*

"With 80% of the US microwave link business concentrated with a small handful of customers, it was decided that the analyzer could best be introduced by having a design engineer from HP Ltd. give the initial customer demonstrations. Feeling was that many of your customers will have special application requirements and we could answer these on the spot.

"Joining me on the demonstration tour will be Finlay Mackenzie who also came to HP Ltd. from Ferranti. Finlay has been the second "man-in-charge" and will be in a strong position to answer any questions that either you or your customers may have. The price of the MLA is $7200 ex. Palo Alto and delivery from production will begin in January. Limited field exposure with the MLA has been very gratifying. As one customer put it, "This is the best instrument HP has introduced since its 8551B Spectrum Analyzer". We hope he is right, and that our sales will reflect his confidence."

The US tour took place from 2 to the 20th October 1967, and the itinerary took in many of the major manufacturers and operators, as well as Bell Labs.

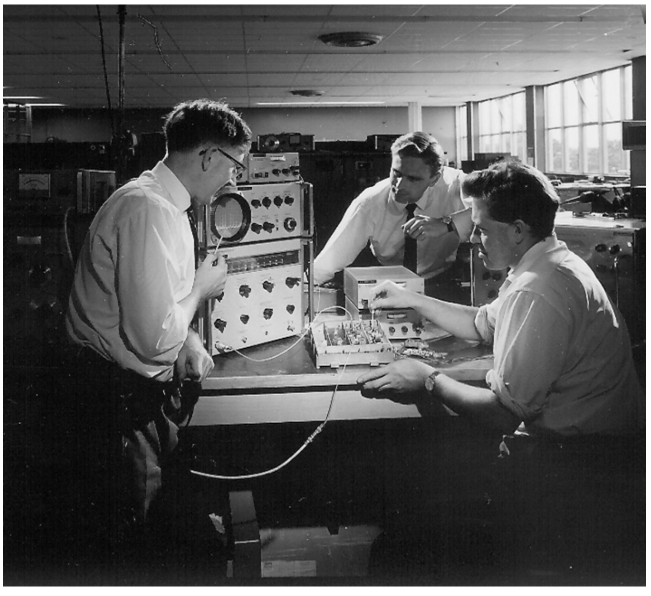

Finlay Mackenzie (left) and Peter Carmichael testing the MLA, with David Leahy watching

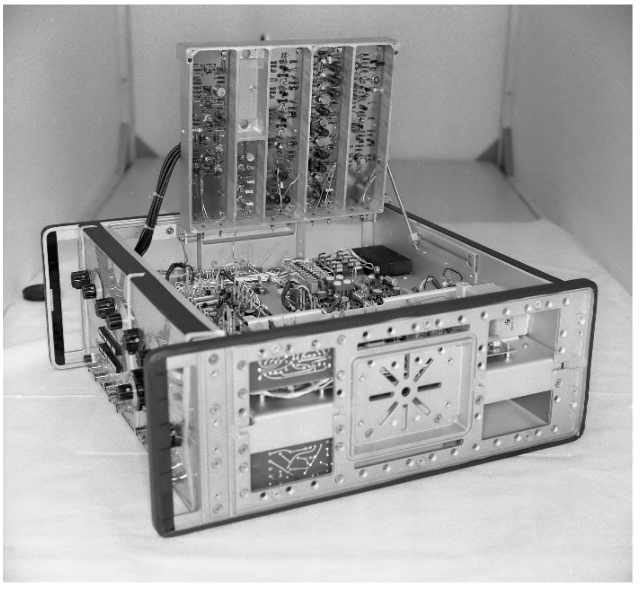

Interior of 3701A MLA showing the RF casting, a similar casting was used in the 3702A receiver

In October 1967, the South Queensferry team got the seal of approval from Corporate when Dave Packard presented a book to the library with the inscription,

"*From the gang at Palo Alto in recognition for the outstanding performance in developing the Microwave Link Analyzer*".

By December 1967, the new MLA (HP 3701A, 3702A, 3703A) was in production and ready to ship the first unit to the Marconi Company. Apparently, quite late in the day there had been a name change as the preliminary data sheet of April 1967 describes the test set as a Communications System Analyzer (CSA), but sometime between then and December it became the more specific Microwave Link Analyzer. Interestingly, the system now preserved in the National Museum of Scotland (T.2010.65.1-3) does include a very early HP 3701A transmitter from mid-1967 which has CSA engraved on its front panel[2].

The MLA and its design team received another accolade on 21st April 1969, when the Division was presented with the Queen's Award to Industry for technical innovation by the Lord Lieutenant who visited the factory. In his letter to all staff, Dennis Taylor, Managing Director, said,

"*I would like to congratulate all of you, and particularly those concerned with MLA project, for their efforts. This really establishes our Company in the UK as a significant contributor to British industry and I am sure that by your continued diligent effort HP Ltd will go from strength to strength.*"

The original sales forecast of 15 units per month was quickly revised as sales went to 30 units or more a month. In fact an average production rate of 30 to 40 units a month became the norm throughout much of the life of the MLA product line over the next 15 years. By the spring of 1971, 1000 units had been shipped and to mark the occasion, the 1000th MLA was handed over to Mr Alberto Fronza from Siemens of Italy in the presence of Secretary of State George Younger MP.

Such was the success of the product in North America, that in the late 1960s a special version was produced for this market (known as the Bell Telephone market). As Finlay Mackenzie recalled, the project nearly didn't go ahead because the corporate managers advised against making something exclusively for "Ma Bell". Fortunately, Finlay persevered despite being told "he'd be better making toasters"! The most noticeable difference in the Bell MLA was the special WECO[3] connectors on the front panel which replaced the standard BNC connectors. The WECOs were large cylindrical gold-plated jacks that gave the instrument a distinctive appearance, and in some cases two connectors replaced one BNC to provide the balanced connection required in the Bell System. Along with some other specification changes to suit the market including

[2] The unit has serial number U728-00051 which indicates it was made in the UK around week 28 of 1967 and with unit number 51 indicating it is a production prototype, so one of the first telecom products produced at the factory. Production units started at 100. The companion HP 3702A receiver has serial number U825-00281 indicating it was made a year later and was the 181st production unit.

[3] Western Electric Company which was the manufacturing arm of the Bell Telephone System

compliance with WECO manufacturing standards, the revised model became the 3701Z/02Z/03Z to distinguish it from the original A-version. It acquired the Bell System tester number KS-20548 meaning it was an approved product so the orders came flooding in from "across the pond". In the factory it was known as the "01Z", the "Bell MLA" and the "KS MLA". Between 1970 and 1973, some 720 units of the Bell MLA were built and shipped to North America. It was the ubiquitous test tool for the world's largest microwave radio network.

Quite a culture built up around the MLA in the factory. The instruments were all hand-wired, and the production staff took great pride in dressing the cable harnesses and connections for a neat appearance. The test engineers developed a certain "black art" in aligning the units for best performance. Although the MLA design was inherently high accuracy, it was always possible to get a bit more by "tweaking". Some of the adjustments were slightly interactive so it was a question of knowing the trade-offs. There were push button attenuators at input and output to adjust signal levels. These were built on special printed circuit boards, and the test engineers would take slivers off the tracks with a scalpel to get the best impedance match. Carve a bit too much off though, and you were in trouble. It was a sort of craft skill, a million miles from electronics in the 21st century. The attenuators were also sold separately as the 3750A and several thousand were manufactured over a 20 year period.

Despite its success, the new MLA was dogged with some serious reliability problems in the early days, as recalled by one of the design engineers, Reid Urquhart:

"*The unreliability had two main contributors. The infamous reed relays, used liberally throughout the MLA system for switching and calibration, and the electrical connections to the PC boards made through crimped connectors which pushed on to pins soldered to the boards. One reed relay in particular used for the group delay calibration, failed a lot since it had to "buzz" at quite a high rate to provide the split trace calibration on the MLA display[4]. I have memories of Ken McDougall spending much of his time travelling around the Globe with his trusty soldering iron and his pockets stuffed with replacement relays, fixing failed MLAs. The troublesome push-on connectors were replaced by soldered joints or by using square pins on the boards to create a tighter fit. I remember a story, current at the time, that Peter Carmichael, the MLA's inventor, was banished to Quality Assurance and told to fix the problem 'or else', which he did and got praise for his efforts. The irony of this was not lost on many at the Division, as Peter was curing the very problems he had designed-in as the R&D project leader!*"

During the life of a product, many changes and improvements were made to enhance performance, manufacturability and reliability. Once the product transferred out of the R&D design phase, these technical improvements were the responsibility of production staff and the production engineering department. In Appendix A3, one of the MLA test engineers, Bill Lauchlan, records his memories of working on the MLA systems in the early 1970s and describes how the engineering improvements were made and sometimes incorporated in earlier customer units. The first MLA was a fantastic product and between 1967 and 1973 shipped over 1200 units in addition to the 720 Bell units – over 1900 in total with revenue of over $14M, which was a lot of money in the early 1970s. A great start for the new Division.

[4] Reid designed an ingenious circuit using diode switching to provide an accurate electronically switched group-delay calibration, which completely eliminated the reliability problem.

Further Developments - The 3710A MLA (Unit Preserved)

Soon after the first MLA began shipping, the R&D team started working on its replacement. No doubt fired-up with the success of the product and ideas on how to make it even better, an enhanced version was planned. The theory of microwave radio links had developed around that time and engineers understood more about the complex impairments that could occur which degraded the traffic signal. In particular, the use of higher-frequency test-tones (ten times higher than those used on the first MLA) began to be a requirement. Finlay Mackenzie recalls the excellent relationship South Queensferry had with Farinon and Collins Radio in Dallas. Their engineers provided valuable inputs for the new MLA. An engineer, M. Ramadan at Collins, had published papers on the use of high-frequency test tones and declared the HP MLA was useless without them! "*We took his input and included the extra test-tones in the new MLA, although many customers were not that interested.*" The new instrument added this capability and enhanced the overall specification. Work started on what was to be known as the 3710A MLA.

It did pretty much the same job as the 3701A system, but with design improvements had even lower residual errors. By early 1972, the new MLA was ready to go into production and over the next 18 months the first MLA was gradually phased-out and the 3710A MLA system took over. The transmitter (the lower unit in this photo) was a significantly new design which included a plug-in unit. It was notable in that it used an arrangement of gears and indicator ribbons for setting and displaying the three key transmitter parameters. It looked nice and gave a setting resolution similar to that of a modern digital numeric display, but the mechanics of the gear-deck were problematic in the early days until all the tolerancing had been worked out properly. Eventually production and the gear-deck ran smoothly. The 3702B receiver looked similar to its predecessor, but with more ergonomic controls and a more modern rectangular CRT replacing the round one of the first model. The new product was described in an article by Reid Urquhart, the Project Leader, in the Hewlett-Packard Journal of September 1972[5].

In May 1972, the South Queensferry division hosted some unusual visitors from Hungary. They came from a telecommunications research lab (TKI in Budapest) and had been working on the theory linking MLA measurements to the equivalent noise levels experienced by the telephone user. As mentioned earlier, cumulative impairments in the radio link create background noise for telephony. The relationship is quite complex and can only be fully evaluated using high-frequency test tones. TKI were interested to collaborate with HP to develop these ideas using the brand new 3710A MLA system.

Tom Sakarny and his two colleagues (plus their obligatory Communist Party minders) visited the factory for a week to make some measurements. The visit was interesting for us too, as TKI brought with them a Wandel und Goltermann ZFM-70 which was the German

[5] HP Journal, September 1972 pp 8 – 16. http://www.hparchive.com/Journals/HPJ-1972-09.pdf

company's competitor to the MLA[6]. Going back to the measurement methods discussed earlier, it took the "gold standard" approach and was basically a very good radio link packaged as an instrument. It was nothing like as accurate as the HP MLA, but it operated like a real microwave radio so was useful for doing the experiments. Tom Sakarny and his colleagues set themselves up in the board room overlooking the Forth Bridge and got to work.

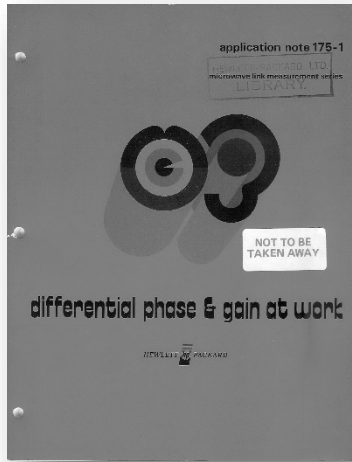

The result was Application Note 175-1 "*Differential Phase and Gain at Work*". It became famous, mainly because hardly anybody could understand what it was all about. Nevertheless, we distributed loads of copies over the years. It made the division look knowledgeable and having a copy lying around was similar to pretending you understood the Theory of Special Relativity!

The 3710A, like many of the second generation products developed by the Division, was a huge seller. It was in production for 13 years and when it was phased out in 1985, 3600 units had been produced bringing in over $50M.

The 3790A and 3711A Microwave Link Analyzer Systems

The next MLA development in 1974 was in response to a new generation of high-capacity microwave links which used a 140 MHz intermediate frequency rather than the standard 70 MHz. The result was the 3790A MLA system. It was a design challenge as the residual specifications needed to be very tight but at the higher frequency it was harder to achieve. The design team used a technology called thin-film hybrids. These miniaturised circuits used an alumina substrate rather than a printed circuit board and were fabricated using chip components, all done in the South Queensferry factory's Thin Film Laboratory. They had inherently wider bandwidth and delivered the required specification. It was described in the HP Journal, November 1975[7], the main author being Ian Matthews, the Project Leader.

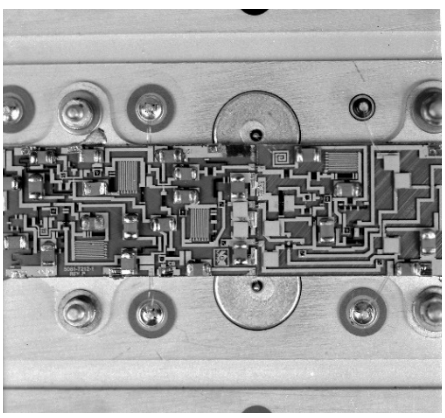

Thin-film hybrid circuit from 3790A

Development work on the 3792A 140MHz IF MLA receiver

[6] Wandel und Goltermann was one of South Queensferry's key competitors in the telecom test business over the years. WuG had owned this market for many years, so must have felt threatened when HP entered the business in the 1970s.
[7] HP Journal, November 1975 pp 13 -24. http://www.hparchive.com/Journals/HPJ-1975-11.pdf

Introduced in 1975, the 3790A system turned out to be a poor seller and only around 70 units were manufactured. It was quickly superseded by a dual IF version covering both the 70 and 140 MHz applications in one unit, and again used the thin film technology. It was introduced in 1979, the transmitter being the 3711A and the receiver the 3712A. In appearance it was similar to the 3790A and indeed the earlier 3710A system.

The analogue measurement performance of 3711A system (shown here) was superb and the accuracy and flatness of the MLA family had gradually improved to the point that it was as good as it would ever need to be, even over very demanding routes such as satellite links. However, by the early 1980s the appearance and user interface consisting of knobs and rotary switches was decidedly “old hat”. By then all the more advanced HP instruments had push-button keyboards, microprocessors and numeric digital displays. In 1975, South Queensferry itself had introduced the company’s first microprocessor controlled instrument, the 3745 Selective Level Measuring Set, so the expertise was certainly there.

Around 1981 work started on a new project, the Third Generation MLA, model 3795A. It would use a new platform developed in California called the Modular Signal Analyzer (MSA) or Modular Measurement System (MMS), which provided the display, operating system and power supply. All that was necessary was to put the MLA measurement hardware in plug-in modules.

Shortly after design work started, a major hiccup occurred when a Japanese competitor, Anritsu, introduced a direct competitor to the 3711A MLA, the ME538 with similar measurement performance, some degree of push button control and digital displays, and worse still a lower price. It was what the 3711A system should have been, but the Division had fallen into the trap of incremental engineering on a successful product thinking it would give the lowest investment for the best return. You can do this for a while, but eventually you end up with an outdated product in the market and one which is probably expensive to manufacture as it doesn’t use the latest technology and production techniques. This was the problem with the 3711A: it used hand-wiring and had to be tested manually.

Division management then wondered if the new 3GMLA could ever make any money in a declining market and with low cost competition. By the early 1980s, the telecom market was going rapidly digital and the analogue microwave link was reaching the end of the road. There was digital microwave radio, but the traditional MLA was overkill for this new application.

Another factor preyed on their minds. In the late 1970s a product had been developed for the radio market called the Baseband Analyzer, a highly integrated and expensive processor-controlled test set (see Chapter 3, page 31). Introduced in 1981, it had very poor market acceptance and was one of the most expensive flops in the Division’s history. Basically it was too late into a market that was in decline, so customers weren’t interested in expensive hi-tech solutions. Their attention was now on emerging digital communications. The fear was that the new MLA would be another Baseband Analyzer.

The project was cancelled. It was disappointing for the project team but allowed valuable design effort to be transferred to new digital test equipment including dedicated test sets for digital radio (see Chapter 4). As for the MLA, HP decided to trade on its long established reputation and push the idea of "*why not stick with the instrument you know and your staff are trained on*". The price of the MLA was dropped about 30%, in the hope that customers would find the decision a "no-brainer". The strategy had some success and took the wind out of Anritsu's sails (or should that be sales). One of the marketing staff, Harry Simpson, came up with the idea that the old MLA had "*mechanical memory*", that is its knobs stayed in the same position when you switched off and would wake up with the same settings when switched on! Who needed a microprocessor?

The MLA struggled on through the 1980s, and in the last few years just the 3711A/12A was manufactured with the last shipment in January 1990. Of this final version, around 1100 units were sold over ten years. So came to an end the 25 year life cycle of South Queensferry's Microwave Link Analyzer family. A total of nearly 7400 units had been sold worldwide, bringing in revenue of over $100M.

Epilogue

It would be hard to overestimate the importance of the MLA in those early years of the Division. Three facets are particularly significant:

- It was a beautifully elegant solution to a measurement problem. Its system design and transistor circuits are still a delight to the electronics engineer over 50 years after they were created. In a way it is a timeless engineering classic much like a fine steam locomotive or an old car.
- It proved there was a valuable market in consulting a target group of customers and designing a dedicated test set for their needs rather than cobbling something together with standard test equipment. This formed the model for the future product development strategy at the Division.
- It was a remarkable business success in the global market and confirmed South Queensferry's position as the specialists for communications test in HP. The revenue it generated in the 1970s enabled the development of many new products in the emerging digital market, through HP's self-financing business model.

Without the MLA the history of HP at South Queensferry would undoubtedly have been completely different.

Acknowledgements

Thanks to Finlay Mackenzie, Bill Lauchlan and Reid Urquhart for help in preparing this chapter

3

Chapter Three

Microwave Radio Products - Analogue

The striking success of the Microwave Link Analyzers, in the late 1960s and early 1970s, convinced the South Queensferry product development and marketing teams to look for further opportunities in the growing microwave radio market. A number of new products came out of this, initially for the analogue microwave radio market and in the 1980s for the new digital microwave application.

In a sense they were all a sequel (chronologically and technically) to the original Microwave Link Analyzer development, and so the next two chapters are also a sequel to the previous one covering the story of the MLA. For that reason, I will not repeat the technical information on microwave links, but in the next chapter will give some detail on how the new digital microwave links differed.

As for the products, they are something of a mixed bag, to say the least. One of the digital microwave products in the 1980s almost rivalled the MLA for market success over its long life, whereas another product turned out to be the biggest commercial disaster in the Division's history!

The 3730A/B Down Converter

In the previous chapter, I described the anatomy of the basic analogue microwave radio, and for clarity the basic structure is illustrated again on the next page. As shown, the MLA was designed to make measurements at the Baseband (BB) or traffic ports on the transmitter and receiver and also the Intermediate Frequency (IF) ports which normally operated at 70 MHz, although also 140 MHz on later systems.

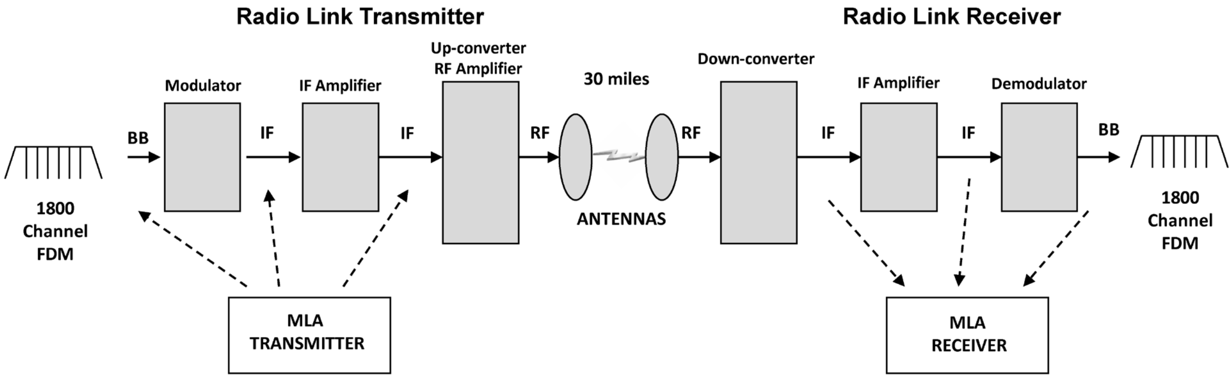

The MLA was particularly useful for measuring the performance of the wideband modulator and demodulator sections and also for characterising IF amplifiers and filters. You could also measure the overall response of the microwave sections (up and down converters, antennas, RF filters and line-of-sight path) by connecting the MLA at the IF interfaces and effectively doing a back-to-back test of the RF section. However, it was not possible to measure the microwave transmitter and receiver separately. The 3730A Down Converter solved part of this problem. The original idea had apparently been suggested by Finlay Mackenzie, one of the MLA pioneers.

It was basically quite a simple product, consisting of a microwave mixer and a tuneable microwave local oscillator. This allowed the RF microwave signal to be down-converted to the 70 MHz IF, suitable for connection to the MLA receiver input, so that the transmitter side of the microwave link could be tested independently.

South Queensferry was able to benefit from development work already done at the Microwave Division in Palo Alto. In the late 1960s, they had developed a number of hybrid microcircuits for their microwave spectrum analyzers and sweepers, including microwave mixers, attenuators and, in particular, YIG tuned oscillators. YIG stands for Yttrium Iron Garnet and it was found that a small sphere of this material would resonate at microwave frequencies rather like an electrical tuned circuit, so it formed the resonant element in the oscillator circuit. The beauty of the YIG was that its frequency could be tuned over quite a large range by altering the magnetic field induced around it. This made it perfect for HP’s swept microwave instruments and also for the down converter[1].

In the case of the 3730A, the YIG was tuned manually until the output IF frequency came in range, and then an Automatic Frequency Control (AFC) loop was switched in to lock the down-converter’s oscillator on to the incoming microwave signal, thus holding the IF centre frequency close to 70 MHz. The AFC loop used a very stable pulse-count discriminator developed to control the sweep linearity and frequency stability of the 3710A MLA transmitter, so that was another module that could be imported into this simple instrument. There was a meter on the 3730A front panel that gave a direct readout of the IF frequency measured by the discriminator, which helped the manual tuning.

[1] The YIG resonator also had a high Q-factor (narrow bandwidth) so this ensured the oscillator had low phase noise and the resulting down-converted IF signal was relatively free from added noise when making a group delay measurement on the MLA receiver.

To cover the range of microwave transmit frequencies found in terrestrial microwave and also satellite links, the 3730A had a range of plug-in units housing different YIG oscillators with an appropriate tuning scale. These were:

- 3736A 1.7 GHz to 4.2 GHz
- 3737A 3.3 GHz to 6.5 GHz
- 3738A 6.3 GHz to 8.5 GHz
- 3739A 10.7 GHz to 11.7 GHz

For receiver testing, Stanford Park Division produced the 8605A Communications Sweep Oscillator as a companion to the 3730A and there was an option so it could be modulated by test signals from the MLA generator. Later, they introduced special plug-in modules for the 8620A Sweeper to do the same thing. There was quite a close tie-up between South Queensferry and Palo Alto on these products.

The 3730A was introduced in 1972 at around the same time as the 3710A MLA system and was described in the September 1972 issue of the HP Journal[2]. The 3730A mainframe was usually sold with one or more of the YIG oscillator plug-ins, so the selling price varied. By 1979, 813 systems had been sold with total revenue of $6.3M, so the average system price over the years was around $7.7k. A total of 995 systems were sold up until 1980 (revenue of $7.5M), when it was superseded by the 3730B.

The second generation "B" model (shown here) was similar in concept, but had a number of significant enhancements. It was intended to partner the new 3711A MLA system introduced in 1979, which had dual 70/140 MHz IF specification. The 3730B also had dual-IF output, with wider measurement bandwidth of 100 MHz on the 140 MHz IF setting. Some new HP YIG oscillators allowed the four plug-ins to span the microwave band from 1.7 GHz to 14.5 GHz, and the tuning frequency was displayed on a digital LED readout. It also had a couple of additional capabilities compared to the first model.

- Firstly, the frequency plug-ins now included the YIG oscillator as well as the microwave mixer and accessories, so the plug-in was a self-contained down-converter with an IF output. This allowed the plug-in to operate outside the instrument on an "umbilical cord", so the microwave connection could be made close to the unit under test, reducing errors due to microwave cabling.

- Secondly, the AFC tracking loop was enhanced. The original unit had a very low bandwidth AFC control loop that only responded to very slow changes in the average frequency, but did not track the low-frequency sinusoidal sweep from the MLA transmitter (typically 18 or 70 Hz). The new model could do that, but it could also be switched to a wider-bandwidth tracking mode, capable of following the low frequency sweep modulation from the MLA. Much as the original MLA receiver had done, the

[2] HP Journal, September 1972 pp 17 – 18. http://www.hparchive.com/Journals/HPJ-1972-09.pdf

tracking minimised the effects of residual non-flatness in the 3730B and improved measurement accuracy. The tracking mode also allowed much wider sweeps, of up 250 MHz bandwidth at RF, to be displayed on the MLA receiver. This was useful for checking antenna systems, waveguides and satellite transponders.

An informative paper on applications of the 3730B including the tracking facility, was written by south Queensferry engineers, Bob Easson and Robin Sharp and was later included in the Telecommunications Measurements book[3].

The 3730B system was introduced in 1980 and was in production for 11 years, with last shipments in November 1991. The table below shows the prices and volumes for the various modules in the system:

Model	US List Price (1986)	Unit Volume	Revenue
3730B Mainframe	$4180	1243	$5.2M
3736B 1.7 to 4.2 GHz	$5360	474	$2.5M
3737B 3.7 to 8.5 GHz	$6430	756	$4.9M
3738B 5.6 to 11.7 GHz	$7500	479	$3.6M
3739B 10.7 to 14.5 GHz	$10715	364	$3.9M

Total sales were around $20M for the complete system, with sales fairly evenly distributed between North America, Europe and International.

Both the 3730A and 3730B were successful products, with one system being sold for every three MLAs. An average of 1.6 plug-ins were sold per mainframe.

3707A Baseband and Sweep Generator

As mentioned above, special options of the 8620 Sweep Oscillator could be modulated with the Sweep and Baseband test signals from the 3710A Microwave Link Analyzer. If the user just wanted to generate the swept RF signal, then the MLA's IF capability was redundant. The idea of the 3707A BB + Sweep Generator was to provide a low cost solution by repackaging the sweep and baseband generator circuits into a compact unit.

The 3707A was launched in May 1980 and cost $1750. It was not a success, and it only sold 28 units before being discontinued in 1983. Although the idea was OK, by then thousands of MLAs had been shipped so the likelihood was that anyone who wanted to do this test would be able to use an existing MLA system. Here it is shown (top) with the 8620A sweeper, on top of a 3712A MLA receiver.

[3] "Enhanced Microwave Radio Measurements using a Tracking Downconvertor" by Robert Easson and Robin Sharp, in Telecommunications Measurements, Analysis and Instrumentation by Kamilo Feher, Chapter 10, pp. 345-360 ISBN 0-13-902404-2 1987

The 3724A Baseband Analyzer (BBA)

Now we come to the most notable instrument in this chapter – the Baseband Analyzer or BBA for short. In the late 1970s it was a *cause célèbre* at the Division but soon became, in some people's eyes, the *bête noir*, as we shall see!

The baseband input and output of the microwave link, as shown in the simple diagram at the beginning of the chapter, were the traffic ports on the system where the Frequency Division Multiplex (FDM) signal was connected to the link. As explained in Chapter 5 on Selective Level Measuring Sets, this is a complex multi-channel analogue signal, typically comprising 1800 independent telephone channels[4]. The microwave link had to transfer this complex signal while minimising the degradation and interference between the multitude of telephone signals. This was why the Microwave Link Analyzer was needed for aligning the various sections in the link for minimum distortion.

Link misalignment caused intermodulation distortion on the baseband signal. With this multi-channel signal, the independent telephone channels were uncorrelated so the composite signal resembled wideband random noise. The effect of the link distortions was mainly to increase the apparent background noise on an individual telephone channel. This varied depending on how much traffic was being carried and where in the FDM frequency spectrum the particular channel was located. Older readers will remember when making telephone calls in the analogue era, how the background noise rose when you made a long-distance telephone call and the connection was made.

This system noise performance was a critical parameter for a microwave link, and it was usual to test it as an overall quality measure after all the MLA alignment tests had been done. The test simulated the effect of loading all but one traffic channel with speech, then measuring the noise level in the one empty channel at the far end of the system under test.

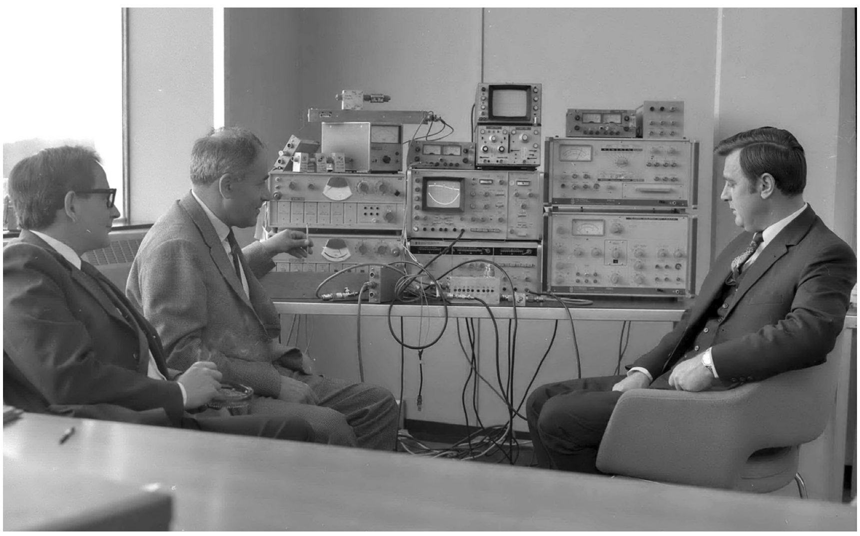

The relationship between MLA measurements and the corresponding noise level in a telephone channel was the purpose of the visit by the Hungarians from TKI Budapest in May 1972. In this photo, Tom Sakarny is adjusting an IF filter using the new 3710 MLA.

[4] Most high-capacity radio links operated at 1800 channels (8.2 MHz bandwidth) using 30 MHz RF channel spacing typically at 4 or 6 GHz, although there were also lower capacity systems with fewer channels. The 140 MHz IF, 11 GHz radios could carry capacities of 2700 channels, with a 40 MHz RF channel spacing.

Next, they would measure the effect on channel noise using two instrument systems from competitors. On the right is the Wandel und Goltermann ZFM-70, a competitor to the MLA, but differing in that it used a very good radio modulator and demodulator as a measurement standard (the "gold standard" approach described in Chapter 2 page 18). This meant it could also be used to simulate a real microwave radio in the noise test. The other test system on the right, is the Marconi White Noise Test Set (TF2091/92) used to measure the noise in the telephone channel, described next. TKI had developed the mathematical theory and wanted to test it with practical measurements, but clearly HP had some gaps in its portfolio of test solutions, so competitors had to be used.

To make this system noise test out-of-service, the radio link was loaded with wideband noise of similar bandwidth and power to the full telephone traffic load of 1800 channels, or whatever, with the empty transmit channel being created by switching in a deep and narrow notch filter at various frequencies in the transmitted band of noise. At the far end of the link, the noise level was measured using narrow bandpass filters corresponding to the frequencies of the notch filters, with the notch filters switched in and then out. The difference between these two readings gave the effective signal-to-noise (S/N) ratio in a traffic channel across the link.

This is called the White Noise Test. The better the link alignment and lower the distortion, the better the S/N ratio. The residual noise at the bottom of the notch would be equivalent to the background noise experienced in a telephone channel in traffic.

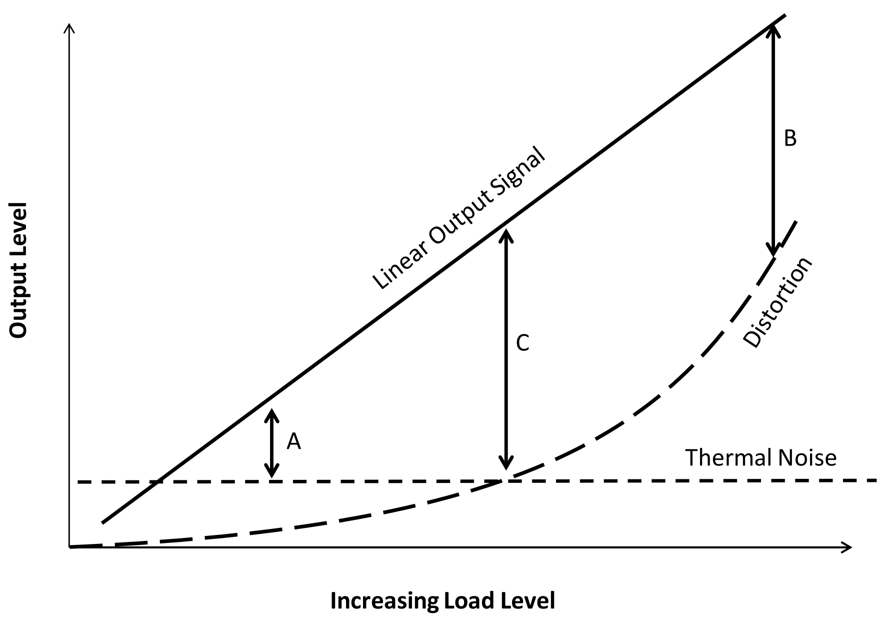

This diagram shows how the S/N ratio is affected by system loading: The noise at the output of a system is a combination of thermal noise (which is unaffected by drive level), and distortion or intermodulation noise which increases as the drive level or system load increases. At low drive levels (point A), the S/N ratio is dominated by thermal noise. At high drive levels (point B), it is dominated by distortion. At the optimum drive level (point C), the contributions from thermal noise and distortion are equal, and maximum S/N ratio is achieved.

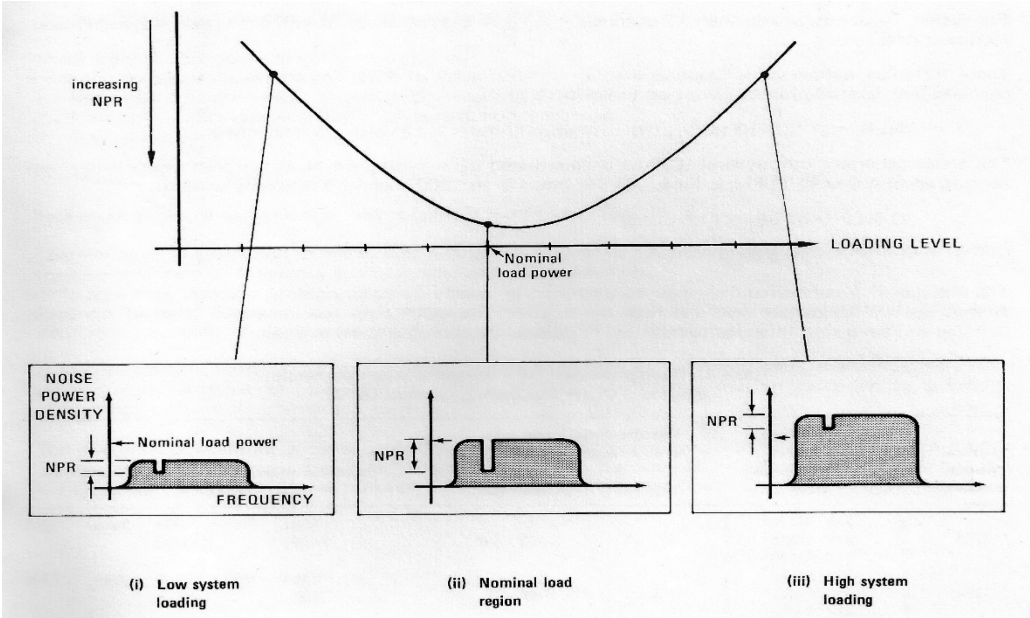

The diagram above shows the key parameters of the white noise test, with the band-limited white noise and notch filter. The usual procedure was to switch out the notch, and take a zero reference level measured through a bandpass filter of the same centre frequency as the notch. Then the notch would be switched in, and the difference would be the Noise Power Ratio (NPR) for that channel frequency, and equivalent to the signal-to-noise ratio on the working system. It was usual to measure at three or four specific notch frequencies across the band of noise. This was defined in an international standard[5]. The diagram also illustrates another facet of the white noise test. If the loading level is too low, then the signal to noise will be degraded by the background noise (thermal noise) of the system, whereas if the loading is too high, the noise will rise due to intermodulation. Between the two is an optimum working level which gives maximum NPR, the bottom of the curve. This graph was referred to as the "V-curve" or "Bucket Curve".

The most popular white noise test set in the 1970s was the Marconi TF2091/92, and it was seen in microwave radio stations as often as HP's Microwave Link Analyzer. It became the industry standard, although there were other products around such as the Wandel und Goltermann RK-5. In 1974, Marconi Instruments published "*The White Noise Book*" by M.J. Tant, cementing their reputation in this business.

The Marconi boxes were quite simple: the transmitter was a wideband noise source along with a set of switchable plug-in filters to define the loading bandwidth as well as the sharp notch filters, while the receiver was a simple wideband level meter preceded by some switchable narrow bandpass filters aligned to the notches. You would take a reference level with the notch switched out, and then, with the notch in, adjust the attenuator to get the same meter reading. The difference in attenuator setting was the NPR value. It was a completely manual test; the clever part was the design of the plug-in filters some of which needed to be crystal-lattice implementation at the higher frequencies.

Buoyed by the success of the MLA in the early 1970s, no doubt South Queensferry wondered if it could also grab a share of the "white noise" business. However, there was

[5] CCIR Recommendation 399-3, "Measurement of Noise using a Continuous Uniform Spectrum Signal on FDM Telephony Radio Relay Systems". The same test was applied to FDM cable systems.

little point in producing a copy of the Marconi set, as HP's offering would probably be more expensive! In any case it was not the HP Way: it had to make a "contribution". Given the Division's interest in digital noise generation, there were some early investigations into using digital processing but this proved impractical. Apparently, it was suggested to Mike Crabtree (the Radio Products' Section Manager in R&D), that a new instrument could do for the range of baseband measurements what the MLA had done for carrier measurements by including all the necessary tests in one unified instrument. Chief among the baseband measurements was the white noise test, but other important testing was done with a Selective Level Measuring Set (SLMS) described in Chapter 5, which could also be used for making flatness measurements in conjunction with a tracking generator. Thus was born the idea of the Baseband Analyzer.

With Queensferry having so much involvement in the microwave radio market through the MLA, there was no shortage of customer inputs for the new Baseband Analyzer (BBA). In fact it could be argued there were too many, as the product definition quickly encompassed just about every measurement that could possibly be needed, regardless of the cost. The inclusion of a baseband spectrum analyzer and sweeper added significantly to the amount of circuitry, and required a good quality CRT display for the results. The list of instruments was to include:

Wideband power meter	Frequency counter
Selective voltmeter (SLMS)	Spectrum analyzer
Synthesized signal generator	White noise test set

Investigations started in early 1974 and development work got underway in 1975 with a team of initially eight engineers under the project leader Richard Roberts. There was a lot of very demanding analogue circuitry to be developed, and such a big multi-function instrument required the largest microprocessor system and software development so far undertaken at South Queensferry.

The overall architecture was monolithic rather than being modular with a high level of integration between the six main measurement functions, reusing various circuit blocks under processor control to provide the instrument features. Although unified, there was so much measurement hardware, the instrument had to be divided into three physical boxes with interconnecting cables so that it had some degree of portability. The main instrument was housed in two standard HP instrument cases 7" (184 mm) high and 20" (508 mm) deep, while the filter modules were housed in one or more mainframes 5.25" (133 mm) high.

The main instrument sections were model number 3724A and 3725A. The 3724A contained the selective receiver and spectrum analyzer hardware, with the main receiver front-end boards housed in a solid milled aluminium structure at the bottom of the instrument (shown here), which was needed to achieve the very low noise floor specified. The receiver back end and measurement circuits were on plug-in cards above, along with the microprocessor. The 3725A contained the CRT (an HP 1340A display module), frequency synthesizer (for receiver local oscillator and tracking generator), and generator output circuits with levelling and attenuation.

Although having different model numbers, the two boxes were completely integrated and could only be operated as a pair, and indeed the operating keyboard was spread across the two front panels[6]. The third box, the 3726A, contained nine plug-in filter modules for white noise testing, and also the noise generator which was fed through the output stages in the 3725A. The 3726A was also controlled by the microprocessor system. More 3726A mainframes could be added if the customer needed more than nine filters. A total of around 30 different filter modules had to be developed, a massive design task in itself.

In terms of circuitry, software and mechanical design it was a monumental project with up to 17 engineers working on it in R&D. There were some outstanding design achievements, not least the design of the selective receiver. Early on, the engineers decided to try and make all selective measurements, including the white noise test, using the wideband tuneable receiver - that is without the inconvenience of pre-selection filters. This was a very challenging objective, since the white noise test involved measuring the residual noise at the bottom of the slot using a 1.74 kHz filter in the presence of the full broadband noise load over a bandwidth of say 12 MHz. The industry specification for a white noise test set was an NPR of 67 dB, which in the presence of the full 12 MHz load, required a receiver dynamic range of nearly 110 dB.

The traditional selective level measuring sets at the time, from Wandel & Goltermann and others, didn't come close. In the early 1970s, South Queensferry developed the ground-breaking 3745 SLMS (as described in Chapter 5) which had a much improved front-end with better dynamic range. It was 10 to 15 dB better than other instruments on the market, but still well short of what was required.

Guy Douglas, who designed much of the BBA's selective receiver, took the 3745 front-end design and improved it further. Unusually for an SLMS, the input attenuator was designed with 2.5 dB steps, rather than the normal 10 dB steps, so that input level could always be adjusted close to the optimum point at the bottom of the NPR V-curve (as shown above). With these enhancements, the new design achieved around 70 dB NPR. Achieving this level of NPR required painstaking optimisation of signal levels at every point in the design, to trade-off thermal noise against distortion[7]. Another critical factor was the need for very low phase-noise in the synthesized local oscillator which tuned the receiver, as this added noise to the measurement and reduced NPR. David Stockton worked on this aspect of the design. There is little doubt that South Queensferry's BBA was one of the best selective receiver designs ever made in terms of accuracy and dynamic range, far outstripping the performance of other SLMS instruments from HP or any other company.

Another significant effort was the microprocessor system and software. As mentioned earlier, it had to control all the various measurement modules, calculate and display results, perform auto-calibration and execute automatic measurement routines. Lawrence Lowe, Brian Woodroffe and the microprocessor team used a Motorola 6800 processor with 56 kilobytes of program memory (ROM), pretty well the maximum that could be addressed. There was also 4 kilobytes of RAM, much of which was used for interfacing to the CRT display. The CRT was used to display the spectrum analysis trace and the swept response for baseband flatness, both of which needed a rapid update rate. To achieve the response

[6] This physical arrangement was similar to HP's 8568A spectrum analyzer of the time, one of the industry's most popular instruments.

[7] Even passive components could not be ignored, for example ferrite-core inductors produce distortion due to magnetic hysteresis.

speed and avoid loading the limited processing power of the 6800, the designers used Direct Memory Access (DMA) allowing the registers for the CRT display to access the RAM on certain phases of the memory cycle[8]. The RAM stored the digital data (8 bits x 256 horizontal points) coming from the analogue measurement circuits.

The CRT was also used to display alphanumeric characters, reading out measurement results and providing menu-driven data entry tailored to the specific measurement function. This was a new idea, and the BBA along with the 3779 Primary Multiplex Analyzer (see Chapter 8) were the first products to use this approach. The BBA took this a stage further by allowing key measurement results to be displayed in large format (16x normal size) for reading at a distance.

There were several fascinating design features in the BBA of interest to the electronics engineer, and fortunately these were described in some detailed articles in the April 1982 issue of the *HP Journal*[9] which featured the Baseband Analyzer on the front cover.

By 1976, the design was progressing well, and the Division Review of September that year predicted customer shipments would start in February 1979 with an average selling price of $20k and potential 5-year sales of $15.2M or 760 units (around 13 units per month).

There was little doubt that the Baseband Analyzer would deliver on its technical performance – the team included some of the Division's best design engineers at the time. Also, considering the size of the project, it was generally well managed in terms of processes and documentation. However, it became clear there was unease amongst some of the Division's management team. They were unhappy about the gradually lengthening development time (it was still a long way from shipping) and the way this massive project was consuming R&D cash and resources.

Towards the end of 1977, Mike Crabtree, who was the Radio Group Section Manager, left HP to take a job with Marconi Instruments in Chelmsford for reasons now unknown. The R&D Manager Bob Coackley, asked me, Hugh Walker, if I would take over as section manager for a combined Analogue Communications Group including both the radio and FDM (SLMS) products. I started this job at the beginning of 1978, and I particularly remember Bob saying that he wanted me to take a hard look at the Baseband Analyzer, as there were concerns about the business prospects for the venture. I wondered if the whole project might be cancelled.

A couple of months later I'd had time to take things in and chat to the design team. I was away on a short trip interviewing undergraduates at university on what was called the "milk round", looking for promising engineers we could invite back to the factory for interview. One evening on my way from London up to Leeds I was waiting for the train to leave the old St. Pancras Station, and wandered up the platform to look at the diesel locomotive. I don't know why, but the idea came up that we needed to split the Baseband Analyzer into a couple of independent boxes (transmitter including noise generator and filters, and SLMS receiver) that could work together as an integrated test system, but be sold separately to give us more marketing options. I remember sketching out some ideas as the train sped north. The CRT display would have to go, which meant we'd lose the

---

[8] The same idea had been used a few years earlier on Queensferry's first microprocessor controlled instrument, the 3745 SLMS.

[9] Hewlett-Packard Journal, April 1982, pp.1-25 http://www.hparchive.com/Journals/HPJ-1982-04.pdf

spectrum analyzer function and the sweeper. On the other hand, we were just about to start development of a second generation SLMS to replace the 3745 (what would become the 3746 as described in Chapter 5), so maybe we could combine the development instead of running two similar projects side by side. One version of the SLMS could major on the FDM measurements and the other on the white noise test application.

On returning to the factory, I told Bob Coackley about the ideas and he said I should go and speak to Peter Carmichael, the Division Manager. Peter seemed to like the proposal, and said with a twinkle in his eye, "*We'll need to send you on some more train trips!*"

However, the product development team and some in marketing were implacably opposed to the whole idea, correctly pointing out it would negate some of the key features of the integrated box. We were also well into the project, so my proposal would require a lot of re-engineering. Status quo won the day, and the "juggernaut" rumbled on to its original destination. Curiously, the venture seemed to gain a life of its own and nobody knew how to stop it. It was a mistake, as the Baseband Analyzer turned out to be nothing but trouble.

At the Division Review in September 1978, the sales forecast of 13/month remained the same as it had done a couple of years earlier although the list price had risen to $25k. However, the ship date had moved out another year to February 1980, and expected development cost was now $1.2M. This was a huge investment at the time as most products took around a third of this: for example the 3762/3763 Pattern Generator and Error Detector had cost $425k.

One problem we found with products having a long development cycle was that they were prone to further marketing inputs along the way, sometimes because the market evolved during the design phase. One of the marketing staff, Bob Hoffman, had worked with HP in the USA and had good contacts with some of the microwave radio companies such as Collins Radio and Farinon. He had lots of ideas for enhancements to the BBA. Finlay Mackenzie, marketing manager, recalled the battle of wills:

"*He always seemed to be coming to me with new additions to the specification. I felt the BBA became a monster because Bob kept adding to the product. I recall a meeting when I had to tell him that the measurement set was too large and out of control, and that the Project Leader, Richard Roberts, was relying too much on his judgement. His response was, 'Finlay, you are good with people and would make a great personnel manager, but you know nothing about product strategy!'* "

At the Division, we had a nickname for this phenomenon – "*Creeping Featurism*". The line usually went that such-and-such a customer would buy 20 units if we included a specific new feature, and the latest inputs always seemed to be the most important. What this did was further lengthen the development time, a sort of vicious circle.

As the BBA evolved, there were some amusing incidents too. As well as its high technical specification, the instrument had some novel automatic measurement features controlled by the microprocessor which allowed it to do end-to-end measurements across a radio link without a control connection. A particularly useful one was the ability to do the white-noise NPR test at multiple slot frequencies without all the switching in and out of filters needed with the Marconi boxes. The noise generator would cycle the noise source on and off with all the required notch filters switched in and the receiver locked on to this at the far end. On the prototype instruments the key for this function was labelled “NOISE ANAL” until it was noted that it could have flatulent connotations. It became “AUTO SEQ” on the final units. Some of the more complex features were a bit obscure even for the design team. I recall one day asking the Project Leader, Richard Roberts, how a particular key/function worked. He said he’d need to refer to the definition document, to which my thought was that if he couldn’t remember, it was unlikely a customer would ever use it. I think the box did become too complicated; it was trying to do too many things.

At the annual Facility Review in September 1979, the BBA business plan showed a higher list price of $32k, while the mature sales still showed at 13 units/month, giving total revenue of $16.6M. By then the expected R&D cost was $1.4M. Also the expected date for first shipment was now November 1980. Every year that went by, it seemed the completion date moved out another six months. It was converging however!

Around this time, more staff came on board. Effort was needed to transfer the BBA to production, developing assembly and test procedures including the writing of automatic production test software for the test racks. Of course, all this activity was more involved than usual because of the complexity of the product. Similarly, technical authors had been working on the various operating and service manuals, the latter being three large volumes with circuit schematics and performance tests allowing repair down to component level.

Doubts were beginning to surface that the 13/month volume might not be achieved. Finlay’s favourite response to dissenting views was the example of the Norwegian PTT and Mr. Moen (not sure how his name was spelt, but it sounded as Mow-en): “*Mr. Moen is going to buy 50 units and he is only 2% of the world market!*”, he would say. Then one morning a bombshell arrived in our in-trays. Long before the Internet, HP had an internal electronic messaging service called Comsys, which batched text messages and sent them overnight. These were printed out by the IT department and distributed in the mail. The message that arrived from the HP sales office in Oslo was that Mr. Moen had jumped ship and joined our rivals Wandel & Goltermann. A few minutes later, Finlay arrived in the office and we told him the bad news, “*Mr. Moen has joined W&G.*” Finlay slumped in his chair and said to the department secretary in his Hebridean accent, “*Monica, get me a whishky.*” Needless to say, Norway didn’t buy any Baseband Analyzers, but unfortunately they were still 2% of the world market. Usually Finlay had an amazing ability to sniff out emerging markets and new product opportunities, however on this occasion it seemed his commercial antennae were a bit wonky.

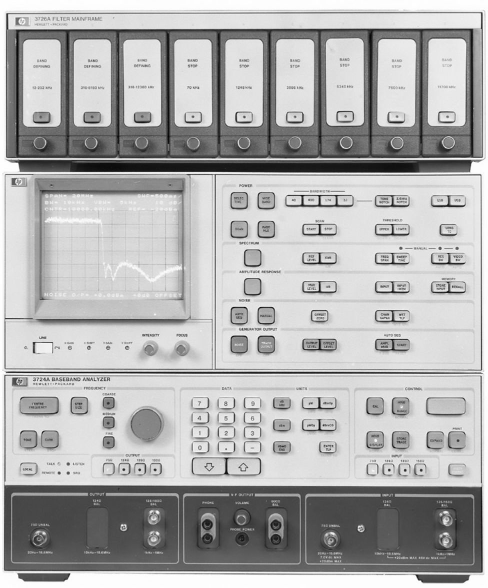

The big day finally arrived in the spring of 1981 when the new Baseband Analyzer was launched in the market. The Product Support Plan, dated May 1981, gave details of some early orders including eight units to South Africa of which more later. The list price was around $40k - $42k depending on the filter modules selected. The plan now showed an expected order rate of 5 units per month, well down on the 13 forecast 18 months earlier. Everyone kept their fingers crossed that it wouldn’t be less than that.

"The White Elephant Test Set"

Sales were a bit slow with 34 systems ordered in the first year, 19 of these coming from South Africa. There was the standard set of promotional literature including a colour brochure produced by Reid Urquhart in marketing which featured a strap-line across the front with the message "*One instrument for all analog microwave radio baseband measurements from 50 Hz to 18.6 MHz*". He was also involved in writing a software package, the 37018A running on the HP85 personal computer, that programmed the BBA to run a sequence of up to 40 measurements and documented the results. Nothing else on the market competed with that.

In 1982, orders amounted to a disappointing 20 systems, as South Africa didn't order anything. The Baseband Analyzer was becoming a bit of a bad joke at the Division and assumed the nickname "*The White Elephant Test Set*", alluding to its white noise capabilities that nobody seemed to want. One problem was that because of the high cost and complexity, very few sales regions would buy a demo unit, so often somebody from the factory had to go out with the instruments (three large transit cases) to do demos – another additional cost. By that time I had moved from R&D to Marketing, and remember on one occasion going to do a demo in Finland. I was so relieved to check the pile of heavy boxes onto the flight at Heathrow Airport, I completely forgot to do the customs paperwork. I was halfway to Helsinki, opened my briefcase and saw the Carnet, the customs document we used for importing and exporting commercial goods without paying duty, which I'd failed to get stamped. Fortunately, I managed to get into Finland and phoned the factory to own up. The reply was, "*Go back to R&D where you belong!*"

We tried quite a lot of promotion. Articles in trade journals and of course the *HP Journal* issue mentioned earlier was an excellent publication to hand out to the sales engineers, customers and indeed anybody with the slightest interest. If anybody bought five units, then it became a "sales success story".

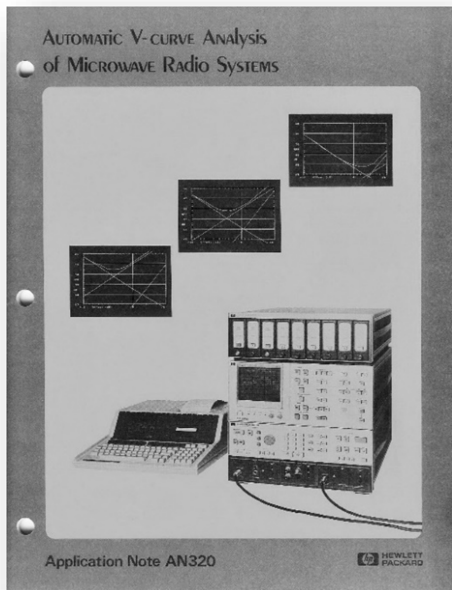

During 1982, the team developed an interesting piece of application software for the BBA and the HP85, which plotted and analysed the NPR V-curves described earlier. This was written up in Application Note AN320 "*Automatic V-curve Analysis of Microwave Radio Systems*". It was a clever idea that related the V-curve back to the measurements made on the Microwave Link Analyzer[10] and a fine article by Guy Douglas (who wrote much of the application software) was later published in the Telecommunications Measurements book[11]. In early 1983 I recall making a trip to Anchorage Alaska with a huge pile of equipment (shipped in advance I'm happy to say) and gave a two-day seminar on the above topic.

---

[10] It related to the theoretical work published in AN175-1 "Differential Phase and Gain at Work", discussed previously in Chapter 2.

[11] "Analog Microwave Diagnostic Measurements" by Guy Douglas, in Telecommunications Measurements, Analysis and Instrumentation by Kamilo Feher and HP Engineers, Chapter 9, pp314-344, ISBN 0-13-902404-2 1987.

Microwave links were important for the oil industry, but it now seems a strange location for the seminar. It was also ice-bound and extremely cold, but the place warmed up a bit in the evening when a multitude of strip clubs opened to entertain the oil workers. The local HP sales guy was an expert, the “Great Alaskan Bush Company” being top of his list. The BBA gave us some tough assignments, but someone had to do it!

The low sales volume created another problem: it was very expensive and inefficient to manufacture such small quantities of a complex instrument. I suspect that the high costs meant that with every unit we shipped, the Division lost more money. The Manufacturing Manager, Alistair Lucas, was constantly lobbying to have the BBA discontinued as it caused so many problems on the production line.

It also had a rather high failure rate, which did not help its reputation with those who bought it. This was mainly due to the large amount of electronics in it, since in general the Division’s product failure rate (measured in %/k$) was better than most. However, the low production numbers meant the population was never really large enough to identify and correct systematic failures. Often it was staff from the Division who had to go out and fix problems in the field. Guy Douglas recalled an exacting trip in the Australian Outback:

“*Six BBAs had been bought to test a new radio link system being installed. It was January with temperatures of 45-50 degrees. I followed the engineers as the units were put in packing cases and thrown into the back of a Toyota 4x4. I got thrown into the back as well and we headed off into the Outback east of Perth. The spare wheel was mounted on the inside of the cabin and held in place by a wing nut. In the course of our 'journey', the wing nut spun its way off the securing bolt and the tyre came away from its fixing. Gives a feel for the kind of vibration and temperatures the instruments had to withstand! In this case I was fitting some special filter boards they needed for their tests.*”

My own BBA reliability story goes back to a trip to China in May 1982, shortly after the HP sales and service office opened in Beijing. I was there for three weeks with two engineers from the Santa Rosa Division and we gave seminars and product demonstrations to the Chinese engineers. The temperatures were high and the air conditioning minimal. One day the BBA, a featured instrument, stopped working as the temperature rose. I took it through to the rudimentary service department and found some sort of intermittent fault in the input circuitry, either a bad connection or sticking relay. I managed to get it going again. It was a complete fluke, but the Chinese were mighty impressed I could give seminars and also fix the instrument!

Sales in 1983 amounted to 36 systems. South Africa bought another 9 and a record number of 13 units were sold in the USA (maybe one in Alaska). In 1984, we finally decided to put the White Elephant out of its misery and gave six months’ notice in case anybody was planning a purchase. Sales that year were 13 units. Some thought the marketing group had given up too easily, but we were flogging a dead horse (or could it be a white elephant).

Total sales were a paltry 103 systems of which 31 were for the South Africa Post Office (SAPO). Total revenue was about $4.2M, but I suspect the Division made little or no profit. Final development costs were around $2.2M, probably over $2.5M once service manuals and transfer to production were included.

The whole venture was a complete disaster and the Division lost several million dollars over a period of nearly 10 years. Fortunately, HP's business model treated R&D as an expense rather than capital investment, so the loss had already been written off against past revenue. At the time, we just forgot about it and got on with some new opportunities, and in those days there were plenty. However, it must have been a bitter disappointment for the staff who worked for many years devising this beautifully engineered system[12]. Looking now at the photographs of the product and reading the *HP Journal* articles, one can't help feeling sad that all that effort went to waste.

One has to admire the perseverance and commitment of the design team who successfully completed this complicated product, despite the cloud hanging over it much of the time. Guy Douglas commented,

"*I have many good memories of working on the team, but as you know, the BBA project was the laughing stock of R&D during much of its development life. And that wasn't easy to live with. Even in the late 1980s and early 1990s, the new management team couldn't resist the temptation to have a cheap laugh by mentioning the Baseband Analyzer.*"

Some of the engineers felt their future careers were tainted by association with the BBA.

The R&D investment that was poured into the Baseband Analyzer in the late 1970s could have funded several alternative new products. Business in the early 1980s was at times a bit flat, partly due to the failure of the BBA. Given the scale of the money involved and the lost opportunity, it must have affected the Division's growth, and probably attracted unwanted attention from Corporate management.

There was fallout from the Baseband Analyzer debacle too. For a start, it put an end to any further developments in the analogue transmission area – the cancellation of the third generation MLA described in the previous chapter, was partly due to this. It was a stark warning of the dangers of a protracted product development cycle, and the Division gradually learned the lesson, culminating in the Fast Cycle Time (FCT) process improvements of the late 1980s and 1990s. When the BBA was devised in 1974, the analogue radio market was buoyant. By 1981 when the product was introduced, the radio manufacturers were moving on to digital systems. The BBA was an excellent unit for production test, but the market had largely disappeared.

For the time, the product concept was wrong. It was an "all or nothing, blot-out-the-sun product", but most customers no longer wanted the full integrated package, preferring to mix and match. There was still business to be done in analogue radio market even in the 1980s. The next product we are going to review, the 3717A 70 MHz Modulator – Demodulator, targeted the same market sector and was introduced about the same time. It was in production for 10 years and sold 850 units, something which speaks for itself.

[12] In 2013 I was in contact with a retired telecommunications engineer, Al Jamison, in Vancouver B.C., who has probably the last working BBA on the planet. He greatly admires the engineering and made the following comment: "Although it may have been a marketing failure, it is an example of the best service diagnostic capability that I have ever seen in any equipment. Whoever developed the maintenance philosophy intended that the unit be able to be serviced for ever – a very thorough design!!"

3717A 70 MHz Modulator – Demodulator

Earlier in the chapter, I noted that one of the competitors for the Microwave Link Analyzer in the early 1970s was a product from Wandel und Goltermann, the ZFM-70 IF Measuring Instrument. As explained, this instrument took a completely different approach to the measurement of modulators and demodulators and the IF sections of the microwave link. Whereas the MLA made all the measurements using specific modulated test signals with a tracking demodulator, the ZFM-70 used the much simpler approach of providing a reference modulator and demodulator that would be used to check the microwave link modules.

For example, the link modulator would be checked against the "gold-standard" demodulator in the ZFM-70, assuming that any measured non-linearity was in the item under test rather than in the test instrument. The same principle applied in reverse for checking the link demodulator. To check IF sections (such as amplifiers, filters and group-delay equalisers) the baseband measurement test sets used the ZFM-70 modulator and demodulator to generate an FM test signal to characterise the IF circuits. However, it was difficult to make the ZFM-70 modulator and demodulator very much better than the units under test, so the measurement errors could be significantly worse than those of the MLA. Consequently, the MLA dominated the market.

Despite these limitations, the ZFM-70 was attractive for some applications because the high-quality wideband FM modulator and demodulator could be used to replace the link mod and demod and operate with wideband traffic signals. It could also be used in white noise test applications (as with the TKI project mentioned earlier) for evaluating the Noise Power Ratio (NPR) with a test set like the Baseband Analyzer. The MLA was incapable of that type of testing since its whole concept avoided the errors inherent in the wideband mod/demod approach by using a tracking demodulator.

Although the white noise NPR test was essentially a baseband measurement, it was sometimes desirable to make the test on an IF/RF link section (for example a series of microwave repeater stations that had no internal modulation and demodulation). This allowed the engineers to estimate the contribution these carrier sections were making to the overall NPR reading for the system.

With the Baseband Analyzer in development, it was logical for the South Queensferry Radio Group to look for a solution to the "gold standard" modulator/demodulator requirement. In the late 1970s, the Division became aware of some excellent modulator and demodulator designs from Northern Telecom in Canada, which they had developed for their 1800 channel analogue radios operating with 70 MHz IF. Whether this was through the HP sales engineers visiting the Northern Telecom customers, or through network operators with experience of the Northern Telecom (NT) radios, is now unknown. The NT modules had a back-to-back NPR of over 60 dB which was 5 to 8 dB better than many other systems on the market. They were good enough to be the basis of an instrument that customers would want to buy and which would rival the WuG ZFM-70.

Some sample modules came in from Northern Telecom, but they were a bit of a mixed bag in terms of performance and there were some concerns about stability and reliability when used in a portable instrument rather than a fixed radio rack. The packaging was ideal however, being a neat and compact aluminium extrusion with connectors at the ends, which would fit perfectly into a standard HP 3.5” (88 mm) high cabinet.

Further sample modules from Canada arrived but these had various problems, and we began to realise the difficulties of dealing with a critical supplier thousands of miles away. It was like “pushing a piece of string”. No doubt staff at Northern Telecom wondered if they wanted to get involved with an exacting customer in Scotland, when all they wanted to do was build microwave links. Some of the project team led by David Haworth, were all for abandoning the Northern Telecom units and designing the modulator and demodulator in-house. Although this was possible, it was unlikely to happen, as the potential revenue from the product was insufficient to justify the extra R&D cost. There was also the risk we might copy bits of the Northern Telecom units having seen them, which would be a problem, although we might have been able to license the design.

Around this time we enlisted the help of a former South Queensferry engineer, Ken Edwards, who had left HP in the early 1970s to join Northern Telecom where he had climbed the “greasy pole”. He helped us speak to more influential people in Montreal, and a better supply contract was duly negotiated following some visits to their factory.
I believe NT eventually set up a separate manufacturing area for the HP modules, where they selected the best performing units and tested them thoroughly. In the end they supplied us with over 900 sets of modules so it was significant business. The three modules we bought from Northern Telecom were a wideband modulator, an IF limiter on the receive side and a wideband demodulator which followed the limiter. This photo shows the three NT modules being tested by Robin Sharp using a computer-controlled BBA.

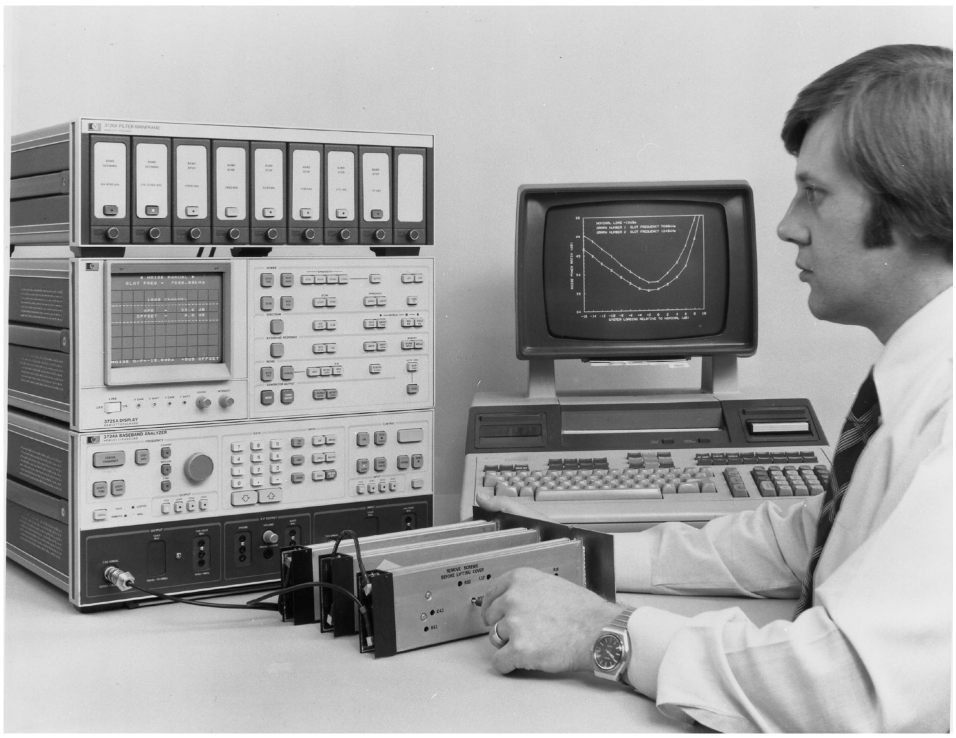

There was still plenty for the design team to do to complete the new instrument that was to become the 3717A 70 MHz Modulator – Demodulator. There were input and output Baseband and IF circuits. On the receive side, the unit had an Automatic Gain Control (AGC) IF amplifier so it could handle a range of input levels. There was an Automatic Frequency Control (AFC) circuit to stabilise the modulator centre frequency.

The NPR of the Northern Telecom units was very good, so the objective was not to degrade the specification significantly with the additional circuitry. The Baseband circuits needed to have very good intrinsic NPR, and the IF sections needed to have a very flat response to minimise carrier section distortions.

Another important development was a range of Pre-Emphasis and De-Emphasis networks to shape the frequency response of the wideband modulating signal. This was standard practice in microwave radio links and the frequency response for different channel capacities was internationally agreed. The emphasis was needed because of a fundamental aspect of Frequency Modulation. For a fixed deviation of the modulated carrier, the power in the sidebands reduces as the modulation frequency increases. With a flat noise floor across the radio link, it meant that higher modulation frequencies were subject to increasingly poor signal-to-noise ratio after demodulation. To overcome this problem, higher frequencies were boosted and lower frequencies attenuated, around a crossover point. At the receive end, this high-frequency lift was compensated by a corresponding de-emphasis network so that the overall frequency response was flat. This meant that the telephone channels across the FDM spectrum would all have similar signal-to-noise ratios. The same principles also applied to transmission of analogue TV video signals. The team designed a total of 16 different emphasis boards covering various channel capacities and TV standards in North America and internationally. Up to five could be optionally configured in the unit and selected by push-buttons on the front panel.

The instrument was completed in late 1980 and shipments began in December, with a forecast of 7 units per month. The basic price was $8100 at launch, but would typically be $9k to $9.5k with emphasis modules fitted. Whereas most of the instrument could be repaired to component level in accordance with HP standard servicing, the Northern Telecom Units had to be replaced by module exchange, so a process for that had to be set up. In the first two years, sales ran pretty close to forecast, but in 1984 and 1985 it rose to 12 units per month. It tailed off to around 5 units/month in the late 1980s, as analogue microwave radio was phased out. The selling price rose considerably in later years.

The 3717A was finally withdrawn in 1991 after total sales of about 850 units. It was a successful product with sales of around $15M. Apart from the white-noise test application discussed earlier, the 3717A was popular as a standalone modulator/demodulator for temporary drop/insert of traffic at repeater stations and for video networking over cable or radio links. Although designed as a companion for the Baseband Analyzer or “White Elephant Test Set”, the relationship ended prematurely!

Some other "Bits and Bobs"

The Microwave Link Analyzers were in production for over 20 years and shipped thousands of units. Along the way, this successful product spawned a number of accessories for various test applications. I’ll conclude this chapter with a brief listing of some of these additional products.

3750A 75 ohm Attenuator

This was simply a repackaging of the 99 dB 75 ohm attenuator used in the first MLA transmitter, the 3701A. It had two rows of pushbuttons that allowed the attenuation to be set in 1 dB steps. It had a top frequency of 100 MHz. HP in the USA made a number of step attenuators but they were mostly 50 ohm impedance, whereas telecom systems always used 75 ohm coaxial interfaces. The 3750A therefore found many applications in the industry. Introduced in 1970, it was in production for 20 years. The initial unit price was $200 rising to $650 in 1990, and 4700 units were sold.

3740A IF Switch

This was a simple solid-state changeover switch that allowed two 75 ohm IF inputs to be connected alternately to a single IF output. It was used with the MLA to measure the differential delay between two paths using the MLA group delay measurement. A typical application was to measure and minimise the differential delay between two receive antennas in a diversity reception system – the DADE test (Diversity Antenna Delay Equalisation). The normal end-to-end group delay measurement would be made across the link and the IF switch operated rapidly to alternate between the two paths. This produced a split trace display on the MLA depending on the delay difference. Introduced in 1973, it cost about $400 but the number sold is unknown.

3743A IF Amplifier

In the early 1970s, the Division established a technology lab to design and fabricate thin-film hybrid circuits. These were for high-frequency applications in the MLA, BER testers, SLMS etc. One of the products was a wideband amplifier with gain of 30 dB, covering the range up to 100 MHz. This was packaged up and sold as the 3743A, mainly for increasing the sensitivity of the MLA receiver. The thin-film circuit had very good amplitude and group delay flatness. It was in production from 1976 to 1985, the selling price started at $700 and doubled over the product life. Unit volumes are unknown.

3744A Baseband Sweeper

This product was designed for use with the MLA to check the baseband flatness of the microwave link up to 15 MHz. The MLA could only check baseband (BB) level at a few points set by the test tone frequencies, whereas the sweeper checked all frequencies and displayed the swept response on the 3702B MLA Receiver. The swept signal was generated by mixing the 70 MHz crystal oscillator output from the 3710A MLA Transmitter with its IF swept signal, offset from centre to sweep from 70 MHz upwards. At the receive end, the level was measured with a diode detector and the DC level displayed on the MLA Receiver. The 3744A also recovered a sweep signal to provide the X-deflection on the MLA Receiver. The product worked well, but was clumsy to use with all the coax cables between the instruments. It might have been better if the measurement had been built into the MLA system itself. The 3744A Baseband Sweeper was introduced in March 1975 and sold for $1500. It was in production until 1981 but only sold around 170 units, a tiny fraction of the MLA sales. It was not really a successful product.

Conclusions

The analogue microwave radio product range, including the MLA family, spanned the range from brilliant success to miserable failure. In the early 1980s, HP had a pretty comprehensive set of test equipment for the design, manufacture and maintenance of analogue radio systems, combining products from Queensferry and the US divisions. This diagram from a publication of the time demonstrates the coverage, and neatly rounds off this chapter before moving on to digital radio in the next.

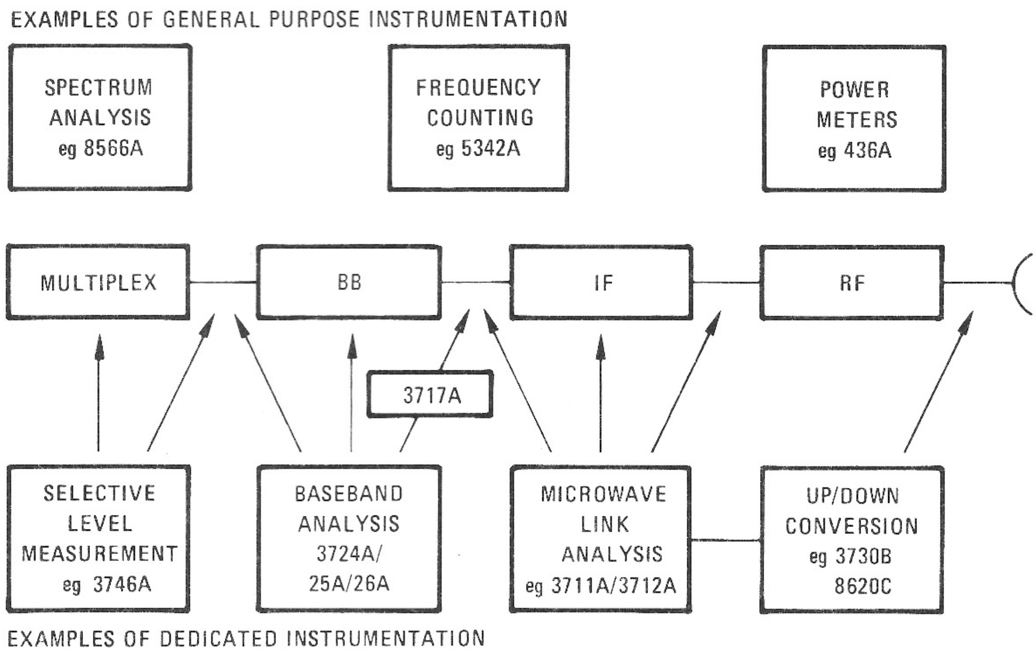

Acknowledgements: Thanks to the following former South Queensferry employees for assistance with this chapter: Boyd Williamson, Finlay Mackenzie and Guy Douglas

4

Chapter Four

Microwave Radio Products – Digital

By the late 1970s, digital transmission was gaining momentum with systems operating over coaxial cable or microwave radio (high capacity fibre optic cable being still nearly a decade in the future). The first digital microwave systems appeared in the mid-1970s, and a popular product in the UK was the early GEC 140 Mb/s system[1] operating in the 11 GHz band (10.7 to 11.7 GHz), which was developed in conjunction with the British Post Office (BT). Here is a simplified block diagram of a digital radio:

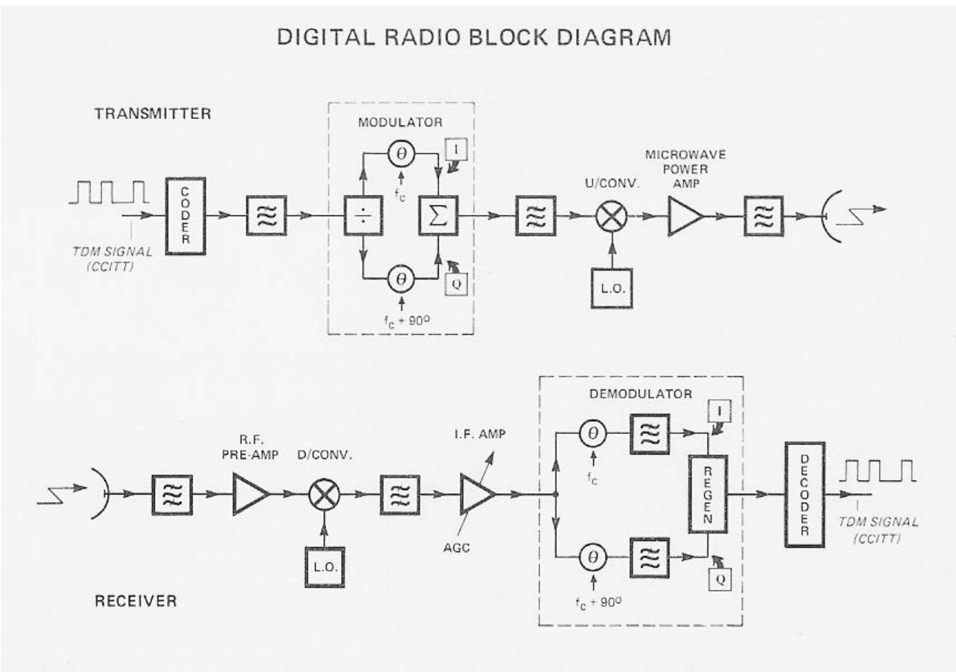

[1] When South Queensferry introduced the 3762A/3763A Pattern Generator and Error Detector at the end of 1977 it had the option of two pattern outputs at 70 Mb/s (half the full rate) specifically to drive the inputs of this digital radio system being installed in the UK.

It is broadly similar to the analogue microwave radio described in the Chapter 2 with the RF and IF sections and operating over a Line of Sight (LOS) radio link, typically 30 to 40 miles in length. The main differences are the modulator and demodulator and of course the traffic input and output which are now serial digital bit streams rather than composite multi-channel FDM telephony.

As shown in the diagram above, the modulator and demodulator are based on quadrature phase modulation. The modulator is supplied by 70 or 140 MHz carrier signals at 90 degrees to each other. They are switched by the digital bit-stream and the outputs combined to form a composite phase modulated signal, which is filtered and up-converted. At the receiver, the phase modulated signal is applied to two quadrature phase detectors (again supplied by two reference carriers at 90 degrees to each other) to recover the digital information from the modulation, which is combined and regenerated to form the outgoing traffic stream. A significant difference with the digital radios was that the receiver had to recover the reference carrier from the input signal and phase-lock to it, otherwise the coherent demodulation would not operate correctly. By comparison, the analogue radio was simpler, since the receiver's frequency discriminator only needed to be tuned to roughly the right IF frequency, not phase-locked to it.

The quadrature modulation produced a family of digital modulated signals, called modulation schemes, depending on how the quadrature modulators were driven by the digital signal. This diagram shows some popular schemes in use in the 1970s and 80s:

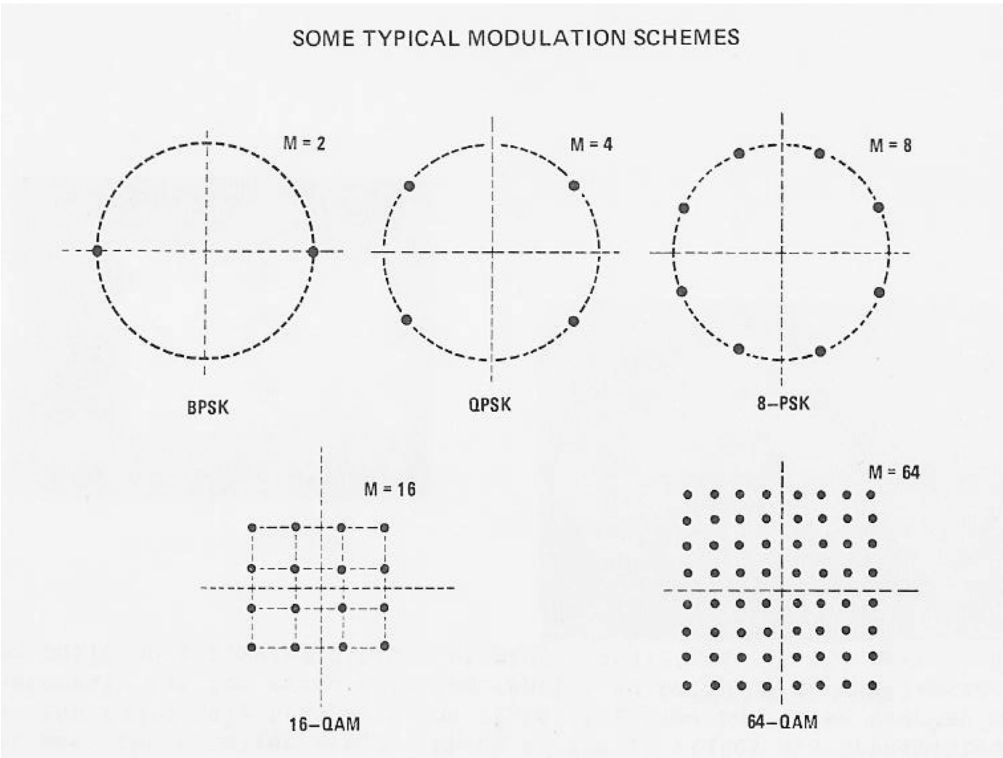

These fell into two categories. The top row all use Phase Shift Keying (PSK) in which the carrier amplitude remains constant (constant radius on a circle) and the phase is switched to different relative positions. The second category, in the bottom row, uses both phase and amplitude modulation and are referred to as Quadrature Amplitude Modulation (QAM) schemes.

Each dot on the diagrams is referred to as a phase state or "symbol", which represents one or more bits of digital information. For the simplest scheme, BPSK or 2-PSK, the two states can only represent either a digital "1" or "0". In the case of QPSK (Quadrature Phase Shift Keying), there are four possible states each of which can represent a unique 2-bit word. Each state of 8-PSK can represents a 3-bit word and each state of 16-QAM, a unique 4-bit word. Thus the incoming digital stream goes through an effective serial to parallel conversion. In the case of 16-QAM, four serial bits are received and encoded as one of the 16 phase states, then the next four bits of data are received and encoded as another symbol, and so on. In the case of QPSK, the symbol rate is one half of the incoming bit rate: for example 140 Mb/s becomes 70 Mega-symbols/second, while 16-QAM quarters the rate to 35 Mega-symbols/second.

The lower the symbol rate, the lower the bandwidth needed to transmit the information; or conversely more data can be transmitted through a given bandwidth, the more complex the modulation scheme. (Readers will note the close similarity between digital radio theory and the design of high-speed voiceband data modems discussed in Chapter 6 on Telephone Line Analyzers.)

The downside of more complex schemes was that they were vulnerable to noise and misalignment since the gap between the states was smaller, so the receiver was more likely to make errors when recovering the signal. The early radios, like the GEC 11 GHz system mentioned at the beginning, used QPSK (2 bits per symbol). QPSK was very robust, but had poor bandwidth efficiency. In the case of the GEC radio, the designers took advantage of the resilience of QPSK and effectively operated two radios simultaneously on the same frequency and same route by using the horizontal and vertical polarisation of the antennas, with the two independent microwave signals operating in orthogonal planes. It would never have worked with analogue radio, but the high immunity of digital systems to noise and interference made it possible. However, it was an expensive solution.

Radio manufacturers quickly evolved their designs to use more complex 16-QAM by the early 1980s, and later 64-QAM and 256-QAM. These radios were much more critical in terms of alignment and so needed more rigorous testing in design, manufacturing and installation. Thus, the need for a new range of dedicated test equipment began to emerge.

The need mostly applied to the performance of the new modulators and demodulators, although with the more complex modulation schemes the carrier sections of the link (the IF and RF parts) were also critical. Some parameters such as flatness across the transmission bandwidth were not dissimilar to the old analogue radios, so the Microwave Link Analyzer had a minor continuing role, however its key measurements on FM modulators and demodulators were now completely redundant.

In the remainder of this chapter, I will look at two highly original products developed at South Queensferry in the early 1980s which focussed on this emerging market.

3708A Noise & Interference Test Set (Unit preserved)

As with all digital transmission systems, the objective of the digital radio was to get the digital data stream across the radio link without any bit errors. Bit errors are usually caused by noise superimposed on the digital signal which occasionally causes a "1" to be interpreted as a "0" and vice versa. In other words, the noise causes the binary signal to be on the wrong side of the decision threshold at the sampling instant as illustrated in this diagram:

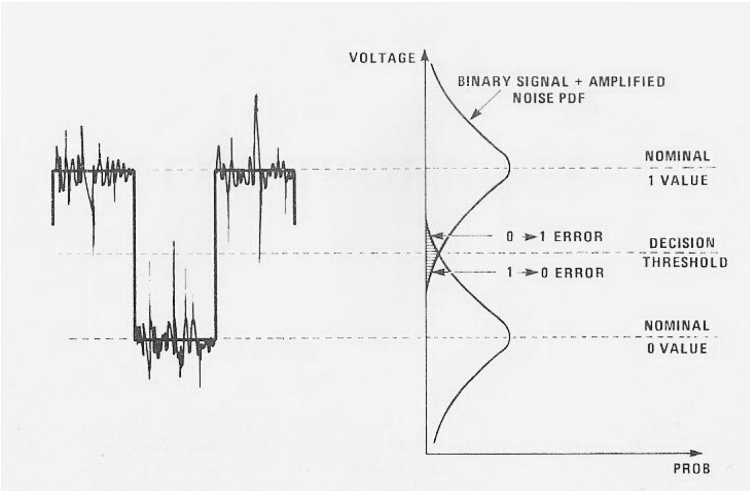

The probability distribution for Gaussian noise is the bell-shaped curve shown above; the two tail areas around the decision threshold line show the probability of binary errors. The higher the noise-level relative to the data signal, the greater the error rate on the system. Superimposed noise can be caused by random thermal noise in the equipment and also by Inter-Symbol Interference (ISI). ISI results from the residue of previously transmitted data adding to the current data bit; in effect there is some short term storage in the system, usually due to finite bandwidth, which causes a tail in the system impulse response. On high-speed data systems such as the high-capacity digital radio, some ISI is unavoidable despite minimisation by proper alignment of the filters and amplifiers in the system. Since the data being transmitted is random or pseudo-random, the resulting ISI has similar statistics to random noise and looks like noise when viewed on an oscilloscope. This composite "noise" appears to reduce the space between the "1" and "0" levels so there is higher probability of error.

The same applies in digital radio, although the signal is more complex. In the 16QAM photo shown below, the 16 phase states of the ideal signal shown as dots on the earlier diagram, become clusters of dots around each mean position, the noise having spread the

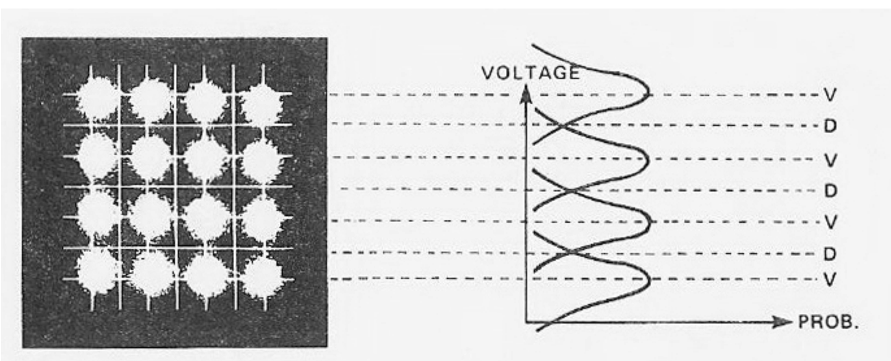

phase state on both axes. Again the effect of noise is to reduce the gap between the states and therefore increase the probability of error.

The relationship between Signal-to-Noise (or Carrier-to-Noise Ratio (C/N) in radio systems) and the error rate, can be derived mathematically assuming perfectly aligned phase states and decision thresholds, and perfect filtering. The mathematics produces a family of curves, sometimes referred to as "waterfall curves", relating the error rate to the C/N for different modulation schemes (shown in the left hand graph). This shows that for a given error rate e.g. 10^{-6}, more complex modulation schemes require better C/N.

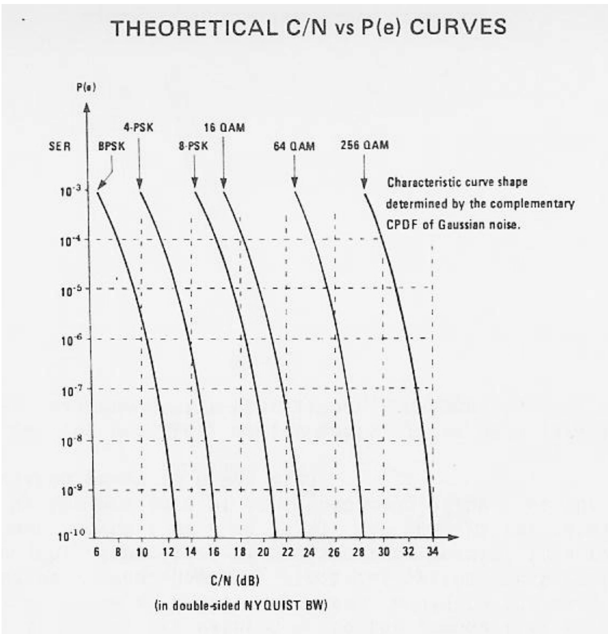

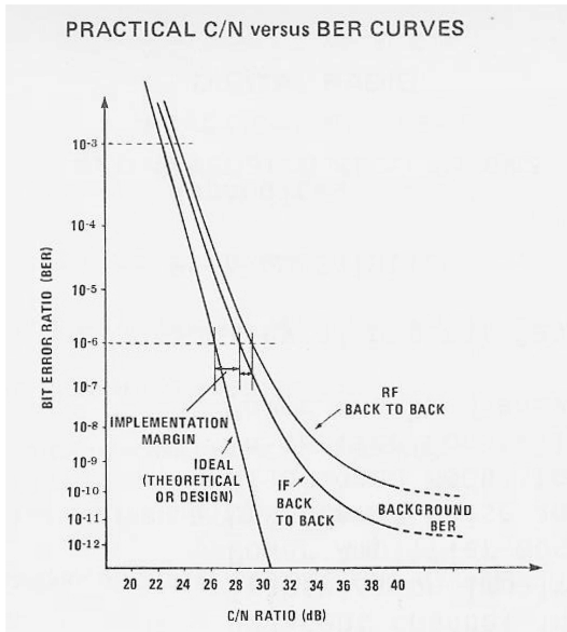

Practical radios can never have better performance than these theoretical curves, and usually the actual performance is measured relative to the theoretical curve to determine the degree of degradation, sometimes referred to as the "implementation margin". Under normal operating conditions, the digital radio receives a strong signal, so the C/N is high and the error rate very low – the background or residual error rate of the radio (typically around 10^{-12}). Under poor propagation conditions (e.g. heavy rain), the receive signal can drop to the level where the error rate rises to a critical level[2].

This allowable drop in received signal from nominal levels to the critical threshold level is called the fade margin. Many factors contribute to the fade margin including transmit power, antennas, receiver noise figure and a whole range of impairments in the radio such as filtering and mod/demod misalignment that cause the implementation margin mentioned above. One way of testing this fade margin in the field or in the factory was to connect a microwave waveguide attenuator at the input of the receiver under normal operating conditions, and determine how much attenuation could be added before the threshold Bit Error Rate (BER) was reached. The drawback with this method was the lack of accuracy and repeatability due to the varying signal levels, particularly across a link. A 1 dB change in C/N results in a 10:1 change in BER. The method was also impractical for measuring the specific implementation margin for modules in the radio such as the modulator and

[2] This is measured as the Bit Error Ratio (BER), usually a threshold of 10^{-6}. The test equipment used is a pattern generator and error detector, such as the Queensferry products in Chapter 9.

demodulator, IF filtering, transmit amplifiers and so on, something which was required during development. A much more accurate and repeatable test method was needed, as the design of radios became more sophisticated.

In the late 1970s, South Queensferry's Marketing Manager, Finlay Mackenzie, became aware of this new customer need. As one of the originators of the MLA, Finlay had an instinctive knack of spotting potential product opportunities, and had heard of interest at Collins Radio in Dallas. He suggested that my marketing colleague, Reid Urquhart, and I should go to Dallas and speak to them. He arranged a meeting first thing one Monday morning, and we planned to fly to Dallas the day before. "*I think you should fly out on the Saturday. This is a really important meeting, you need to be there! What if the flight is cancelled?*"

We ignored his advice and flew down to Gatwick on Sunday morning, to catch a flight to Dallas with an airline called Braniff who operated a Boeing 747. After boarding the flight, we were told we'd have to get off again as there was a problem with the aircraft. Repairs needed parts brought in by helicopter from Heathrow since 747s didn't normally operate from Gatwick. We were on and off the plane twice more and by evening it looked as though the Dallas meeting was in jeopardy. Maybe Finlay had got this prediction right as well! Finally, at about 11 pm, we were all bundled onto the plane once again, the doors slammed and the plane got away as quickly as possible before the airport closed for the night. We arrived in Dallas about 6 o'clock in the morning, with time just to drop our bags at the hotel, get some breakfast and then go to Collins. It must have been a good meeting as we came away with the basic blueprint for the new Carrier to Noise Test Set. We visited other customers too, and the definition of the new instrument gradually came together. It was a brand new idea so there was nothing else on the market to use as a benchmark.

The basis of the product was a calibrated white noise generator which would inject very precise levels of noise into the IF path of the digital radio receiver to create an accurate C/N ratio. A key feature of the new 3708A was an injection assembly which allowed the radio's IF path to be looped through the instrument transparently. The 3708A measured the carrier power of the IF signal and added filtered noise to this signal to create the desired C/N ratio.

The main elements of the design, apart from the injection assembly, were an accurate "true-RMS" power meter with wide range and fast response, and a wideband Gaussian noise source with special band-limiting filters. The specified flat bandwidth of power meter and noise source was 10 MHz to 200 MHz, to cover the main 70 and 140 MHz IF radios and also some lower frequency specialist applications.

As with any noise measurement, the power reading only makes sense in conjunction with the specified noise bandwidth, either that of the noise source or after filtering by the digital radio's IF filters. The 3708A contained several bandpass filters which were individually measured in the factory and an accurate value of the noise bandwidth stored in the specific instrument on the production line. The noise bandwidth of a filter is the equivalent bandwidth of a perfect rectangular filter that would give the same output power when stimulated by the same flat spectrum noise source as the practical filter being measured. Using the noise bandwidth calibration of the filter, the accurate power meter in the 3708A allowed the processor to calculate the "noise density" of the source, which is the equivalent power in a 1 Hz bandwidth (designated by the symbol N_o). With this value, the 3708A could make a variety of different calibrations on injected noise and could also measure the noise bandwidth of other external filters, as Project Leader Geoff Waters recalled:

"*During the development of the 3708, I came up with the idea that it could be used to measure noise bandwidth, a feature that as far as I know remained unique to that product.*"

The product had some other features that enhanced measurement accuracy. Any errors in the fixed attenuators used in the noise source and power meter were stored in software during production testing, so improving the power meter accuracy to around 0.1 dB. Its fast response time (around 10 milliseconds) meant that it could track variations in the carrier power of the received signal via the injection assembly, and the processor automatically adjusted the noise level to keep the C/N ratio constant, greatly improving the stability and repeatability of measurements.

The noise source was designed to have a very flat spectrum over the band up to 200 MHz and also had a very high crest factor (the ratio of the peak value to the mean). Earlier, we saw how the statistics of Gaussian noise created a small probability of errors in a digital signal at the receiver due to the peaks in the noise signal. In order that the injected noise from the 3708A would create the same statistical probability as the theoretical model, its crest factor needed to be very high. The instrument specification was greater than 15 dB, but typically achieved about 25 dB. This required very careful design of the amplifiers in the noise generator circuits to ensure they were operating well below their clipping point.

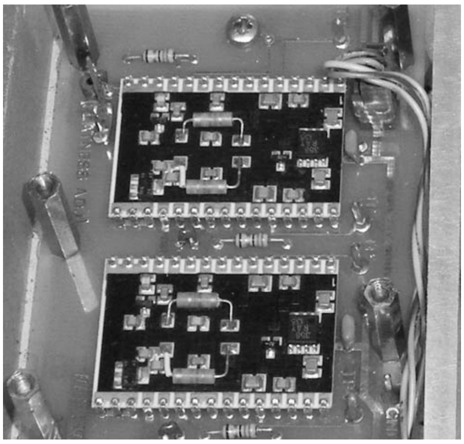

The overall design needed quite a lot of wideband amplification, firstly for the noise generator which amplified thermal noise from a resistor, and secondly for the power meter to amplify signals to a level suitable for the RMS thermal converter which gave the "true" power measurement for signal or noise. The Project Leader, Geoff Waters, came up with the design of a two-stage amplifier (a series-feedback transistor stage followed by a shunt-feedback stage), which had a very stable gain of 10 dB. The design used high-frequency transistors with a cut-off frequency of 5 GHz and the overall frequency response had a flatness of 0.02 dB up to 200 MHz. The circuit was implemented using thick-film technology, and the 10 dB gain module was used in 14 positions within the instrument. Here is a photo of two of the modules mounted in a screened compartment.

The design of these modules and many other interesting features of the instrument are described in considerable detail in a series of papers in the July 1987 issue of the *HP Journal*[3]. In fact the whole issue is devoted exclusively to South Queensferry's digital radio instruments, and has a striking front cover of a radio mast overlaid on a constellation. The papers include a description of the theory and overall design of the 3708A, another on the design of the power meter, one on the design of the noise generator and finally an article on a software product (3708S – 37080A/B/C) which automated the plotting of the C/N versus BER curve.

[3] HP Journal, July 1987, pp19-39 http://www.hparchive.com/Journals/HPJ-1987-07.pdf

An input from customers was they would also like to add an interfering signal as well as, or instead of, filtered white noise, to check how the radio performed in the presence of co-channel or adjacent-channel interference. When used instead of the noise, the 3708's tracking capability could also be used with the interfering signal. Another interesting application of the interference input was to use a non-synchronised sinusoidal interference signal. In contrast to noise, this had the effect of spreading the constellation phase states in a deterministic way compared to random noise – the constellation became a series of circles rather than the cloud of dots round each phase state. The South Queensferry engineers derived some mathematical equations for using this technique to estimate the very low background or residual BER on the radio, thus avoiding the long measurement times normally required to check for errors. This, and many other aspects of digital radio testing, were explored in a comprehensive paper by Geoff Waters (3708 project leader), published in the Telecommunications Measurements[4] book in 1987.

Design work on the 3708A Noise & Interference Test Set started in late 1980, with a design team of around seven engineers. The product was launched in the second half of 1984 and appeared for the first time in the 1985 HP Catalogue. The price at launch was $13.5k, although it rose considerably in later years. Volumes were in the range 12 to 20 units per month.

I recall, however, that the new sales brochure for the launch of the 3708 caused a bit of a hiccup. In the mid 1980s, relations with the field sales force, particularly in the US, were a bit strained at times. The Division Marketing Manager, Mike Cunningham who had come from the USA, was almost paranoid about upsetting the field or creating any risk of ridicule. He advised me that sales literature should use "*earth colours – browns, greens and blues are OK*". The 3708 had been photographed against a purply-pink backdrop for the main instrument shot. When the 10,000 copies arrived from the printer, Mike Cunningham was not impressed. "*We've shot ourselves in the foot goddammit, this looks like the colours of a boudoir. We'll be a laughing stock.*" He was ready to ditch the lot – £8000 worth. We persuaded him against this as it would delay the launch, and anyway being a new product we'd soon need to reprint the brochure, next time using "earth colours". Nobody else was that bothered about it, but some people do get hung-up about such things.

The 3708A was designed for the digital microwave radio market and was particularly attractive for the R&D and manufacturing segments, though it was frequently used in field installation. By the early 1990s, the digital radio market had started to level off, mainly because the much higher bandwidth (ten times higher bit-rate) of the new fibre-optic transmission systems was displacing microwave radio from trunk routes in the network. One might have expected 3708 sales to decline accordingly, however they remained strong throughout the 1990s. This was due to a new emerging application, the market in direct digital broadcast satellites for TV, a huge market requiring the development of satellite receivers and transponders. The Noise & Interference Test Set got a new lease of life.

[4] "Digital Radio Measurement Techniques" by Geoff Waters, in *Telecommunications Measurements, Analysis and Instrumentation* by Kamilo Feher and HP Engineers, Chapter 7, pp180-238, ISBN 0-13-902404-2 1987.

The original design concept and the accuracy of the product were so good that the 3708A never really suffered any significant competition, and became the industry-standard instrument for this type of testing, much as the MLA had done in the 1970s. The 3708A was finally withdrawn in 1999, nearly 15 years after it entered the market, and that was mainly because some of the components it used were becoming obsolete. It was one of Queensferry's longest-lived products and remained largely unchanged during its production life. By 1999, the price had risen from the original $13.5k to an eye-watering $32k!

This table gives an idea of how consistent the sales of the 3708A were over the years, and how its selling price rose:

	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99
Unit Volume	150	139	161	237	259	*180*	*170*	161	150	158	247	138	140	*130*	*50*
Price ($k)	13.5k	13.5k	14.5k	15.1k	15.9k	16.9k	18k	21k	24k	26k	27k	28k	29k	30k	32k
Sales ($M)	2	1.8	2.3	3.5	4.1	3	3	3.4	3.6	4.1	6.6	3.8	4	3.9	1.6

Around 2450 units were produced with total revenue in the region of $50M, very profitable revenue in later years. It was one of Queensferry's most successful products which started with a meeting in Dallas, Texas …. that nearly didn't happen!

The preserved unit is a production prototype (PP003) dating from 1984.

Interior showing hybrid amplifiers and filters

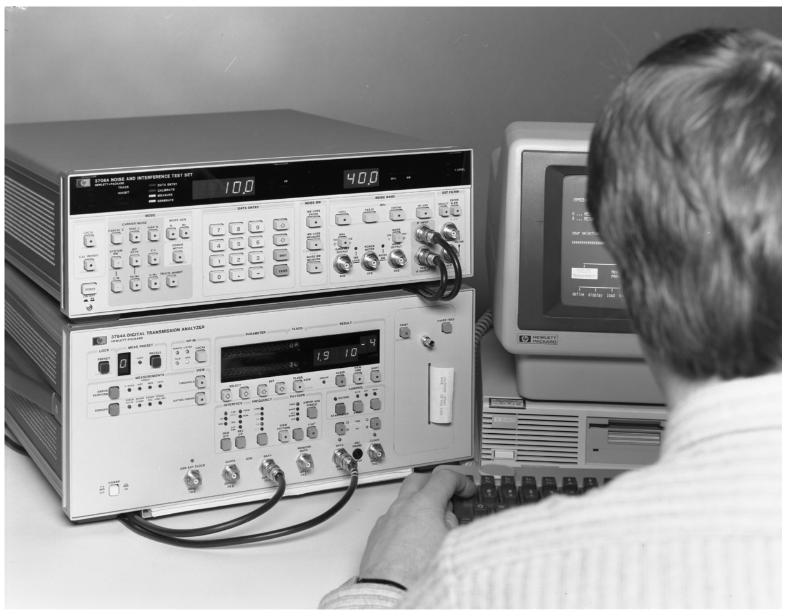

3708A with 3764A Digital Transmission Analyzer under software control plotting BER vs. C/N ratio

3709A/B Constellation Display (Unit Preserved)

The Western Electric (WECO) factory at Merrimac Valley[5] near Boston, Massachusetts, was a major operation from the 1960s to the 1990s, designing and manufacturing much of the transmission and switching equipment used by AT&T across the USA. Merrimac Valley was a very important customer for HP and particularly South Queensferry, and many former engineers from the factory will recall the HP sales engineer, Vince Yaras, who handled the account and had been with HP since the early 1950s. Maybe Vince no longer understood all the latest technology, but he had a huge network of contacts at Merrimac Valley and could usually set up a good meeting for us when visiting from South Queensferry.

I recall a visit to Merrimac Valley in the early 1980s, probably 1981, and as we walked along the corridors of this huge site, I was immediately struck how many of the staff knew Vince Yaras personally and asked about products and orders. Our meeting was with some engineers in the microwave radio lab, presumably to do with the 3708 which was under development. After the meeting, one of the WECO guys said to me, "*Come next door, there's something I'd like to show you.*"

They were developing a 16QAM radio and had the receiver constellation displayed on an old Tektronix sampling oscilloscope, rather like the constellation picture shown a little earlier in this chapter. They used the two channels of the scope to display X versus Y, or in this case the In-Phase (I) versus Quadrature (Q) signals from the phase detector, and had synchronised the timing off the radio's symbol clock. It required some rather complicated adjustments to get a good picture, but they then demonstrated how the constellation displayed various imperfections in the radio, in particular the non-linearity of the transmit power amplifier, a Travelling Wave Tube (TWT). "*Can HP make something that can do the job better?*" And so was born the idea for the 3709 Constellation Display. Geoff Waters remembers around 1983 making further visits to Merrimac Valley, Bell Labs and AT&T and returning with a draft specification for the Constellation Display.

Back at the South Queensferry, Boyd Williamson, who had recently completed the 3746A SLMS, started work on a product definition. The basics were fairly straightforward. What was needed was the front-end of a dual-channel sampling oscilloscope and a timing circuit that would trigger reliably and with very low jitter from the radio's symbol clock, the clock that timed the update of the constellation points. Typically this would be a maximum of 35 to 40 Mega-symbols per second, although the instrument operated up to 80 MHz.

As shown in the diagram below, the Constellation Display would be connected to the In-Phase (I) and Quadrature (Q) monitor points in the radio's phase demodulator. There were two main modes of operation. One was to display the constellation phase states (I versus Q, with I being the horizontal X-direction and Q on the Y-direction). The other was an eye-diagram display of either I or Q versus time, as one might see on a regular sampling oscilloscope. The main difference was that the Constellation Display had an auto-ranging timebase so that it automatically displayed two symbol periods of the incoming radio clock.

[5] At its peak it employed 12000 people on a huge site. It later became Lucent Technologies after AT&T was broken up in 1984. After the Dotcom bubble burst in 2001, it got into trouble and was a victim of the subsequent merger between Alcatel and Lucent, the site finally closing completely in 2007.

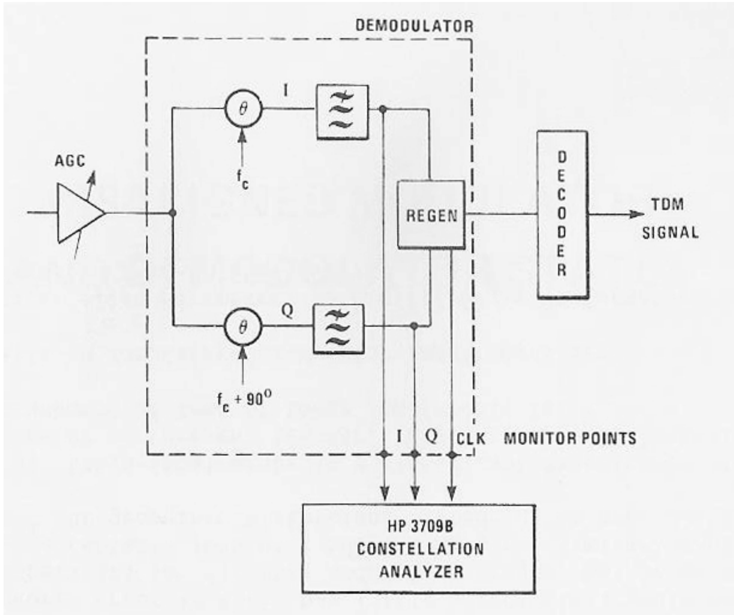

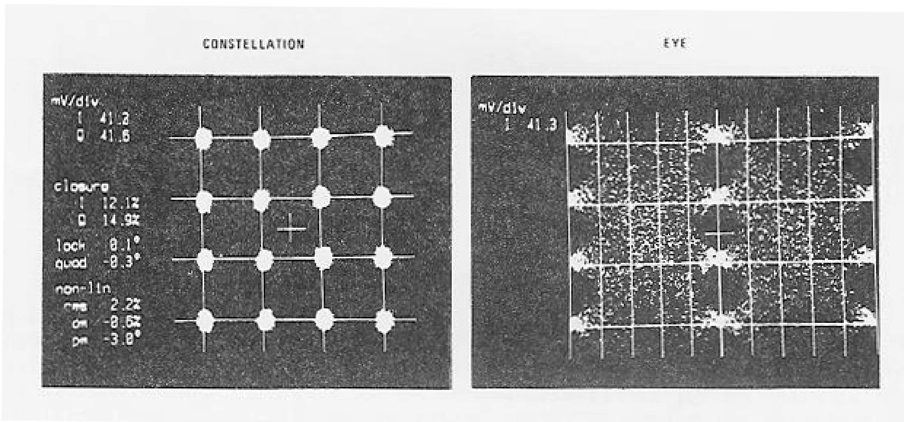

To obtain the constellation display on the left, the sampling instant is precisely aligned with the point of maximum eye-opening in the centre of the screen on the right. Normally this would be the same instant on both I and Q channels. To minimise any error from the residue of previous samples, the sample-and-hold circuit was reset to zero before the next sample was taken. At the sampling instant, a diode-ring was switched on for about one nanosecond (1ns) to record the instantaneous voltage. This photo shows one of the sampler boards from the 3709. The record of hundreds of samples was gradually built up on the CRT screen (a 1340A display module) which used a fairly long persistence phosphor, as was the practice with analogue sampling oscilloscopes. The 3709A had better performance than a normal sampling 'scope as it was optimised for this particular application and it was easier to use. The engineer could do as the lab engineers at WECO had done, and use the display to evaluate and optimise the radio qualitatively. However the other big contribution of the 3709 was the numerical analysis of the constellation parameters displayed at the left side of the constellation display shown above.

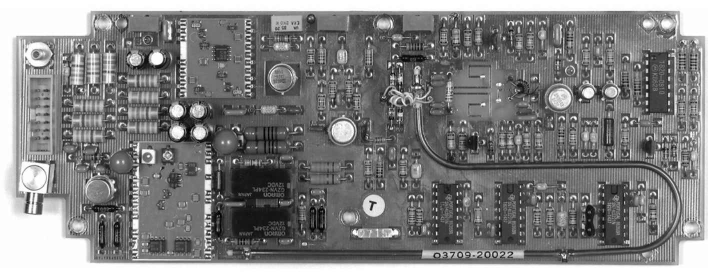

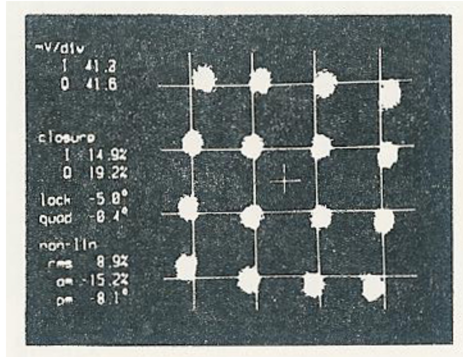

Pressing the MEASURE key on the front panel accumulated 600 samples from the constellation (I and Q values) and converted them to digital data using an Analogue to Digital Converter (ADC). Algorithms in the microprocessor analysed this data and provided statistical averages for several key parameters. These included "Eye Closure" which was a measure of the spread around each ideal constellation point.

Another was the position of the mean constellation points relative to the ideal positions on the graticule which was displayed electronically for the different modulation schemes covered (set by a back-panel switch). From this, the algorithms computed values such as the quadrature misalignment, phase lock, and the level of non-linear distortion in the transmission path. This photo shows the effect of over-driving the transmitter power amplifier which caused both compression of the outer phase states and also some phase rotation.

The design of the Constellation Display was described in some detail in an article the July 1987 issue of the *HP Journal*[6], the same issue with the 3708 articles referred to earlier. A second article explains how the statistical analysis was done to provide the constellation measurements.

After the project started, David Haworth took over from Boyd Williamson as project leader and headed a team of five engineers. The 3709A Constellation Display was introduced at the end of 1985 and started shipping in the spring of 1986. It sold for $8.75k – much cheaper than a sampling 'scope. Initial sales were a bit slow, with 83 units ordered in the first year. Probably this was because it was a new measurement method that was unfamiliar to customers.

About this time, an investigation was started into what further performance data could be derived from the constellation. Geoff Waters, who had been appointed Section Manager for digital radio products, remembered being asked to identify follow-on products for the 3708 and 3709.

"*Boyd Williamson and I were IEEE Com Soc members and sat on the IEEE Radio Committee, the meetings of which coincided with major events such as the International Communications Conference (ICC). This is where the ideas for a future Constellation Analyzer (3719A) came from. We also identified the need for a Multi-path Fade Simulator.*"

---

6 HP Journal, July 1987, pp 4-18 http://www.hparchive.com/Journals/HPJ-1987-07.pdf

Tom Crawford, one of the R&D managers, also held a visiting chair at Edinburgh University with the intent of initiating research into new technology. As we have seen, impairments in radio systems produced geometrical changes in the shape of the constellation, and this suggested an automated pattern recognition system as a diagnostic tool, using the real-time constellation data from the 3709. Tom invited two academics from Edinburgh University (Dr. Peter Grant[7] and Dr Colin Cowan) to join with himself and Virgil Marton to do some research on a machine-learning expert system applied to constellation analysis. It was a reasonable idea, as the principal of deriving transmission impairments by digital processing of the system impulse response had already proved successful with the 4948A In-service TIMS (see Chapter 6 on Telephone Line Analyzers), and digital radio designs were already beginning to use adaptive transversal digital filters to compensate for propagation problems caused by multi-path fading. Nothing commercial came of any of this, however. There was no clear customer interest, it was largely an R&D-driven idea.

In 1987, a "B" version of the Constellation Display superseded the 3709A. It was similar, but had slightly enhanced results display and incorporated measurements on eight different modulation schemes from QPSK up to 256 QAM. There was also a High Impedance Probe Accessory (15709A) which allowed the instrument to connect to radios not having monitor points. In 1989 the price had risen to 11.2k and $13k by 1991. Sales of the 3709B peaked in 1990 with sales of 140 units and gradually declined in the early 1990s. It was discontinued in 1994 by which time the price had risen to $16.9k. Combined sales for the 3709A/B were about 800 units, with revenues of $9M.

It was a reasonably successful product, however its market life and sales volume were limited by several of factors. Firstly, not all digital radios had the required I and Q monitor points necessary for making the measurement – even after lobbying manufacturers and operators. Secondly, as already mentioned, high performance digital radios increasingly used digital adaptive filters (which followed phase detection) to compensate multi-path fading and other imperfections, so the analogue outputs of the I and Q phase detectors only told part of the story of overall link performance.

Interest in the constellation measurement gradually declined as a result. Finally, the dominance of high-capacity fibre optic transmission in the early 1990s displaced microwave radio from trunk routes and it became more of a niche technology for feeder connections, private networks and routes over difficult terrain.

The preserved unit is a 3709A and is a pre-production prototype from 1985.

---

[7] Professor Peter Grant was later head of the Engineering Department at Edinburgh and had a particular interest in communications, digital signal processing and adaptive filters. See http://www.ee.ed.ac.uk/~pmg/cv.html . Research projects he was involved in with HP are cited:
W K Wong "Techniques for Adaptive Multipath Compensation in Microwave Digital Radio Systems", (part-time studies sponsored by Hewlett-Packard) 1987. K E Brown "Expert Systems for Telecoms Fault Diagnosis" (SERC CASE Award with Hewlett-Packard, South Queensferry), 1988. K.E. Brown, C.F.N. Cowan, T.M. Crawford and P.M. Grant, "The Application of Knowledge-Based Systems for Fault Diagnosis in Microwave Radio Relay Equipments", in IEEE International Communications Conference (ICC) Proceedings, pp. 1159-1163, June 1987.

The Digital Radio Business

After the 3709B was launched in 1987, there was no further digital radio development at South Queensferry. Thereafter, the Division's interests were focused on digital transmission analyzers for the new optical SONET/SDH networks in the 1990s, and also Gigabit High-speed BER testers for optical components. (See Chapter 10).

The "baton" for digital radio passed to the Stanford Park Division (SPD) in Palo Alto, California[8]. In the late 1980s they introduced a vector signal generator and vector analyzer (8780A/8980A), high-performance laboratory instruments for doing the I/Q measurements, with a price tag to match. In 1991, they introduced a novel product for simulating multi-path fading (11757A) which could be used to check the effectiveness of the radio's adaptive equalizers. Multi-path fading occurred when more than one signal was received, usually due to atmospheric refraction under certain weather conditions. Apart from the direct line-of-sight path, secondary slightly delayed signals interfered with the main signal and created slopes and notches in the radio bandwidth. This had a severe effect on the demodulated signal. Had not the charter for this area been transferred to SPD, South Queensferry may well have developed the product. Geoff Waters recalled that in the mid-1980s they had worked on the definition of a multi-path fade simulator with BT Research Labs. "*They had patented a novel approach, and we actually paid them to develop a prototype.*"

During the 1980s, HP dominated the digital radio test market. Apart from the novel products from South Queensferry and Stanford Park, HP's popular spectrum analyzers were programmed with the regulatory mask that needed to be used to check the radio's spectrum occupancy. We ran seminars on the theory and measurements. I compiled a couple of slide presentations and remember in 1985 the UK field sales organisation arranged a week of seminars in HP's Central London office. These were very successful with customers coming considerable distances for the full-day seminars. Tom White (UK District Manager) and the team of communications sales specialists had done a great job of promoting the seminars, and interestingly a few years later Tom became the Queensferry Telecom Division's Marketing Manager.

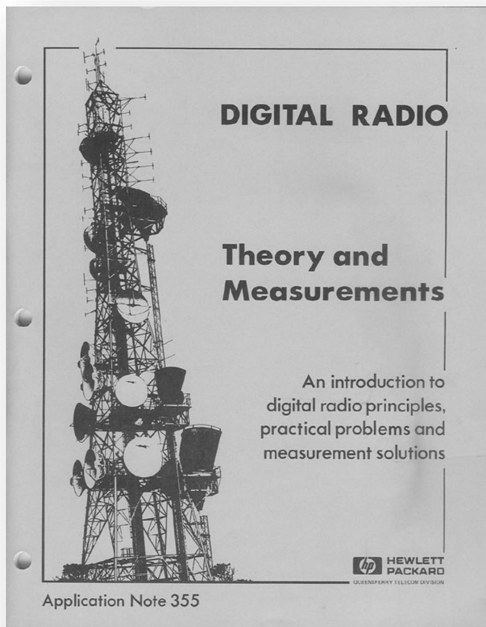

In 1987, I combined the two slide presentations into a single 88-page publication, Application Note 355, "*Digital Radio: Theory and Measurements*". The diagrams in this chapter are taken from this booklet. We used this all over the world, and I remember giving seminars in South Africa and the USA, and Boyd Williamson recalls taking the seminar on the road across the USA and into Hong Kong, New Zealand and Australia. Later, Geoff Waters used a

[8] Although the Telecom Division at South Queensferry was no longer involved in digital radio after 1987, the Queensferry Microwave Division as a development partner to the US microwave divisions had an interest in Digital Radio Test Systems in the 1990s.

much expanded version of this Application Note to teach the M.Sc. course on digital radio for the University of Limerick (Ireland).

Finally, in the early 1980s, an academic, Dr. Kamilo Feher from Ottawa University in Canada, latched onto the digital radio work going on at South Queensferry. He had published a textbook in 1981, "*Digital Communications, Microwave Applications*". He became friendly with a number of engineers at the Division, as Geoff Waters remembered:

"*Boyd and I knew Kamilo Feher very well – we even went to his house in Canada for dinner and swam in his pool! For many years, I considered him a genuine friend, but I lost contact after I stopped attending ICC (International Communications Conference).*"

In 1987 Feher collaborated with Geoff Waters and South Queensferry engineers on a book "*Telecommunications Measurements, Analysis and Instrumentation*", which has been referenced frequently in these chapters.

Acknowledgements

Thanks to the following former South Queensferry staff who have contributed to this chapter: Boyd Williamson, Tom Crawford and Geoff Waters.

5

Chapter Five

Selective Level Measuring Sets and FDM Surveillance

From the historical perspective, there were three products developed in the late 1960s and early 1970s that had a profound effect on the future direction of the South Queensferry division.

The first was, of course, the Microwave Link Analyzer, described in Chapter 2. It demonstrated the value in developing an instrument tailored to the needs of a specific group of customers in the communications market. It was a very successful design which achieved excellent market acceptance, confirming South Queensferry's future as the centre of expertise for telecom test. The revenue stream it generated helped to fund a number of new products in the 1970s.

The second was the 3760/61 Pattern Generator and Error Detector, an early entry into an emerging market of digital communications in the 1970s. It led to a whole family of products in digital transmission (see Chapters 9 and 10), stretching 35 years into the future, ultimately forming the bedrock of the Division's business.

The third product was the 3745A/B Selective Level Measuring Set (SLMS). It was a highly innovative product in its own right, but it also led to two further developments which steered the Division in new directions, perhaps unforeseen at the time. This chapter describes this remarkable instrument and the ramifications that followed.

Frequency Division Multiplexing

Frequency Division Multiplexing (FDM) goes back a long way in telephony, to the beginnings of electronics and the vacuum tube or thermionic valve. Before that, longer distance trunk telephone calls had to be routed through individual pairs of open wires strung along telephone poles at the side of a road or railway line, each pair carrying only one long-distance phone call.

By the early 1930s, a system had been invented to send 12 telephone calls simultaneously through one pair of wires using FDM. The basic 3.1 kHz telephone channels[1] were up-converted[2] to one of 12 frequency slots in the band 60 kHz to 108 kHz, with channels spaced at 4 kHz. This was called a Basic Group, as shown in the diagram below, and allowed 12 telephone channels to be carried over an open wire system that previously carried just one analogue telephone channel.

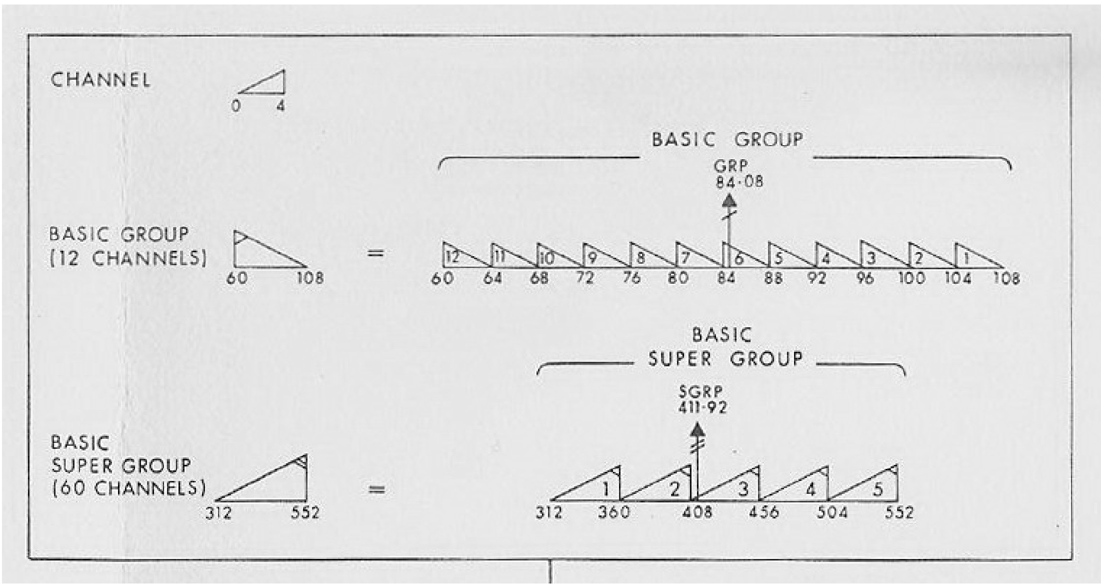

An early example of this in the USA was the J-Carrier System introduced in 1934 and along with later versions such as the N-Carrier System, was quickly established for long-haul interstate telephony. Every six miles along the open wire system, a vacuum-tube repeater amplifier boosted the signal so that considerable distances could be covered without loss of signal[3].

In the late 1920s, engineers at Bell Labs optimised coaxial cable design[4] for low-loss transmission of high-frequency signals, and the same group at Bell Labs patented the crystal lattice filter, necessary for the precise and stable filtering required in FDM systems. By the late 1930s the Bell Labs team was designing a new high-capacity FDM system using coaxial cable. This was designated the L-1 System carrying 600 bidirectional telephone channels through a pair of coaxial cables, with first installation in 1941. The

[1] In carrier systems (either analogue or digital) the telephone channel bandwidth was defined as 300 Hz to 3400 Hz since this was found to give acceptable speech quality. In FDM with 4 kHz channel spacing, this allowed a “guard band” for rejecting the adjacent channel with filtering.
[2] Effectively single-sideband suppressed-carrier
[3] Apparently, a Bell Labs engineer, Harold Black, who had been wrestling with the problem of reducing distortion in repeater amplifiers, invented the concept of negative feedback while travelling on a Manhattan ferry in 1927. It was a brilliant invention which not only reduced distortion and improved frequency response, it also stabilised gain against component aging and temperature changes. In fact it is hard to imagine these long-distance FDM systems would have worked reliably without it.
[4] Originally patented in 1880 by the British electrical engineer Oliver Heaviside, the Bell Labs engineers found the design impedance for lowest loss in air-spaced coax was around 75 ohms.

system had a top frequency of 3 MHz, but the low losses in the coaxial cable meant that repeater amplifiers were only needed every 8 miles.

In 1956, the first Trans-Atlantic Telephone cable, TAT-1, was brought into service. It was a momentous occasion and a triumph of FDM engineering with over 50 submerged vacuum-tube repeater amplifiers in each direction on the deep ocean floor at 37 mile intervals. These were powered through the cable from high-voltage supplies at the shore-ends. The system provided 36 telephone channels (three Groups), although this was later increased to 48 channels by reducing the channel-spacing from the standard 4 kHz to 3 kHz.[5] TAT-1 was a joint venture between the British Post Office and Bell Labs, Bell supplying the repeater amplifiers and the cable being laid by “Monarch” the Post Office’s cable ship, the only one in the world that could lay such a long cable in one go. TAT-1 carried the “hot-line” between Washington and Moscow during the Cold War. The UK landing point was just south of Oban, and a connection was made to the International Exchange in London via Glasgow using a new 900 channel FDM coax system.

Over the years, the capacity of FDM submarine cables increased dramatically and long-distance systems were installed in the Pacific and Indian Oceans. One of the last analogue systems was the Trans-Atlantic TAT-7 installed in the early 1980s which provided over 4000 channels. It used 660 submerged transistor repeater amplifiers. These later systems used a single coaxial cable with band-splitting to provide bi-directional traffic. Submarine systems had a long service life: TAT-1 was decommissioned in 1978, 22 years after it went into traffic.

Further developments of the L-carrier systems in the USA continued until the L-4 system was introduced in the late 1960s. By this time the repeater amplifiers were designed with transistors and spaced every 2 miles with a capacity of 3600 telephone channels per coax and a top frequency of 18 MHz. The L-4 cable had 10 coax tubes in each direction, with one reserved for protection, so the overall route capacity was more than 32,000 telephone channels. This became the main backbone of the North American telephone network. A final development was the L-5 system introduced in the mid-1970s. This was a 60 MHz system with a repeater amplifier every mile, capable of over 10,000 channels per coax and route capacity of more than 100,000 channels.

The same evolution took place in Europe, with the favoured capacity being 1800 channels (8 MHz with 4 mile repeater spacing) and the 2700 channel 12 MHz systems with a 3 mile spacing. Some 60 MHz systems were also installed in the late 1970s.

These high-capacity FDM systems used several standardised methods of hierarchical multiplexing. Starting with the Basic Group of 12 channels described above, five of these Groups were multiplexed (or frequency shifted) to create a Basic Supergroup of 60 channels in the frequency range 312 kHz to 552 kHz. From this 60-channel unit, a number of higher-level structures or frequency plans were created. In the UK, 15 Supergroups were multiplexed to give a 900 channel Hypergroup with a top frequency of 4 MHz. Three Hypergroups were multiplexed to give the 2700 channel 12 MHz system, as shown on the next page.

[5] This became common practice on FDM submarine cables as it gave a valuable increase in capacity with only slight degradation in voice quality, since channel bandwidth was reduced to 2.5 kHz. Another technique used to increase capacity was Time Assigned Speech Interpolation (TASI) which doubled the effective capacity by switching traffic between channels during gaps in speech.

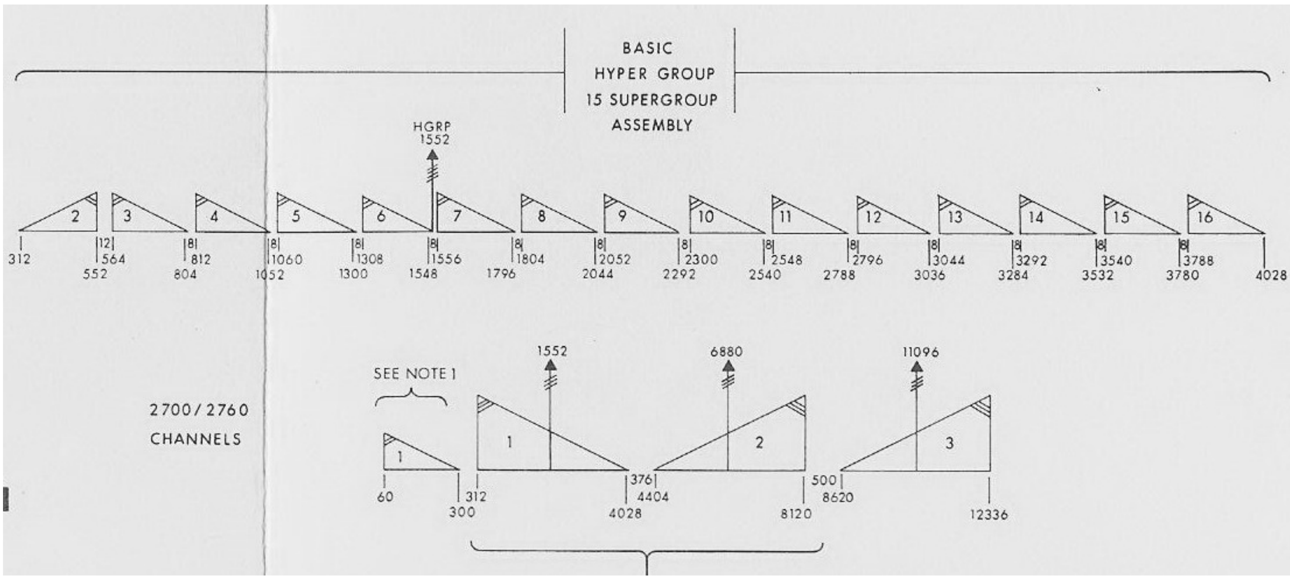

In North America, 10 Supergroups were multiplexed to give a 600 channel Mastergroup with a top frequency of around 3 MHz, and six Mastergroups could be multiplexed together for a 3600 channel 18 MHz system such as the Bell System L-4. A number of variations of these schemes evolved over the years in both Europe and America, and became "standardised" with their own specific frequency plan. Each channel would have a specific 4 kHz slot in the FDM plan.

Although now completely obsolete in the digital age, these FDM transmission systems were once the backbone of the telecom network and the foundation of the enormous worldwide expansion of telephony from the 1950s to the 1980s.

Operation and Measurements

Being a completely analogue multiplexing and transmission system, FDM suffered from accumulating degradations depending on the length of the system. Noise, distortion and gain variations could all cause problems unless carefully managed. In particular, the overall level of the multi-channel FDM signal needed to be carefully controlled to guarantee the best performance. If the level was too low, then background (thermal) noise becomes a problem. If levels were too high, the many amplifiers in the system could start to overload and create distortion. With a multi-channel signal, this also appeared as increased background noise. Thus levels needed to be maintained at an optimum point.

There were many repeater amplifiers along the line, a mile or two apart in high-capacity systems. These were needed to boost the signal, but also had to compensate for the significant roll-off, or increased signal loss at higher frequencies, in the coax cable. When installed, the system would have good equality of level (flatness) across the FDM spectrum, but aging and temperature effects could easily change things, leading to non-optimum transmission levels and degraded performance.

Since the telephone traffic is continuously varying, it could not be used to monitor system gain and operating levels. Instead, special pilot tones were inserted in the multiplex structure to monitor levels and control system gain. For example, each group of 12 channels had a Group Reference Pilot (GRP), usually inserted mid-band between channels

6 and 7. Similarly, a reference pilot was inserted in the Supergroup, Mastergroup, Hypergroup and so on. A high capacity system could therefore have several hundred pilot tones used for monitoring levels and the automatic control of system gain. In a well aligned system, the pilot tone levels would be within +/- 1dB of the correct level.

Even in well-regulated FDM systems, level problems could still exist if some telephone users transmitted too high a level in their particular channel. This was a problem with the growing use of data modems in the 1970s. These "hot-channels" could sometimes cause overloads on the system and increase noise for other system users.

Maintaining the FDM transmission network involved routine measurement of the various pilot tones and also searching for high-level telephone channels, usually in response to complaints from customers.

TEST TONE AND GROUP REGULATING PILOT FREQUENCIES FOR CCITT AND WECO L-600 600-CHANNEL MODULATION PLANS

(TEST TONE FREQUENCIES DERIVED FROM 1000 Hz: GRP REG PLT IS 104.08 kHz OR 84.08 kHz AT BASIC GRP)

CHAN #	BASIC GROUP	GROUP #	BASIC SG	CCITT SG1	CCITT SG2	CCITT SG3	CCITT SG4	CCITT SG5	CCITT SG6	CCITT SG7	CCITT SG8	CCITT SG9	CCITT SG10	WECO L-600 SG9	WECO L-600 SG10
1	107	1	313	299	313	803	1051	1299	1547	1795	2043	2291	2539	2173	2787
PILOT	104.08		315.92	296.08	315.92	800.08	1048.08	1296.08	1544.08	1792.08	2040.08	2288.08	2536.08	2175.92	2784.08
2	103		317	295	317	799	1047	1295	1543	1791	2039	2287	2535	2177	2783
3	99		321	291	321	795	1043	1291	1539	1787	2035	2283	2531	2181	2779
4	95		325	287	325	791	1039	1287	1535	1783	2031	2279	2527	2185	2775
5	91		329	283	329	787	1035	1283	1531	1779	2027	2275	2523	2189	2771
6	87		333	279	333	783	1031	1279	1527	1775	2023	2271	2519	2193	2767
PILOT	84.08		335.92	276.08	335.92	780.08	1028.08	1276.08	1524.08	1772.08	2020.08	2268.08	2516.08	2195.92	2764.08
7	83		337	275	337	779	1027	1275	1523	1771	2019	2267	2515	2197	2763
8	79		341	271	341	775	1023	1271	1519	1767	2015	2263	2511	2201	2759
9	75		345	267	345	771	1019	1267	1515	1763	2011	2259	2507	2205	2755
10	71		349	263	349	767	1015	1263	1511	1759	2007	2255	2503	2209	2751
11	67		353	259	353	763	1011	1259	1507	1755	2003	2251	2499	2213	2747
12	63		357	255	357	759	1007	1255	1503	1751	1999	2247	2495	2217	2743
1	107	2	361	251	361	755	1003	1251	1499	1747	1995	2243	2491	2221	2739
PILOT	104.08		363.92	248.08	363.92	752.08	1000.08	1248.08	1496.08	1744.08	1992.08	2240.08	2488.08	2223.92	2736.08
2	103		365	247	365	751	999	1247	1495	1743	1991	2239	2487	2225	2735
3	99		369	243	369	747	995	1243	1491	1739	1987	2235	2483	2229	2731
4	95		373	239	373	743	991	1239	1487	1735	1983	2231	2479	2233	2727
5	91		377	235	377	739	987	1235	1483	1731	1979	2227	2475	2237	2723
6	87		381	231	381	735	983	1231	1479	1727	1975	2223	2471	2241	2719
PILOT	84.08		383.92	228.08	383.92	732.08	980.08	1228.08	1476.08	1724.08	1972.08	2220.08	2468.08	2243.92	2716.08
7	83		385	227	385	731	979	1227	1475	1723	1971	2219	2467	2245	2715
8	79		389	223	389	727	975	1223	1471	1719	1967	2215	2463	2249	2711
9	75		393	219	393	723	971	1219	1467	1715	1963	2211	2459	2253	2707
10	71		397	215	397	719	967	1215	1463	1711	1959	2207	2455	2257	2703
11	67		401	211	401	715	963	1211	1459	1707	1955	2203	2451	2261	2699
12	63		405	207	405	711	959	1207	1455	1703	1951	2199	2447	2265	2695
1	107	3	409	203	409	707	955	1203	1451	1699	1947	2195	2443	2269	2691
PILOT	104.08		411.92	200.08	411.92	704.08	952.08	1200.08	1448.08	1696.08	1944.08	2192.08	2440.08	2271.92	2688.08
2	103		413	199	413	703	951	1199	1447	1695	1943	2191	2439	2273	2687
3	99		417	195	417	699	947	1195	1443	1691	1939	2187	2435	2277	2683
4	95		421	191	421	695	943	1191	1439	1687	1935	2183	2431	2281	2679
5	91		425	187	425	691	939	1187	1435	1683	1931	2179	2427	2285	2675
6	87		429	183	429	687	935	1183	1431	1679	1927	2175	2423	2289	2671
PILOT	84.08		431.92	180.08	431.92	684.08	932.08	1180.08	1428.08	1676.08	1924.08	2172.08	2420.08	2291.92	2668.08
7	83		433	179	433	683	931	1179	1427	1675	1923	2171	2419	2293	2667
8	79		437	175	437	679	927	1175	1423	1671	1919	2167	2415	2297	2663
9	75		441	171	441	675	923	1171	1419	1667	1915	2163	2411	2301	2659
10	71		445	167	445	671	919	1167	1415	1663	1911	2159	2407	2305	2655
11	67		449	163	449	667	915	1163	1411	1659	1907	2155	2403	2309	2651
12	63		453	159	453	663	911	1159	1407	1655	1903	2151	2399	2313	2647
1	107	4	457	155	457	659	907	1155	1403	1651	1899	2147	2395	2317	2643
PILOT	104.08		459.92	152.08	459.92	656.08	904.08	1152.08	1400.08	1648.08	1896.08	2144.08	2392.08	2319.92	2640.08
2	103		461	151	461	655	903	1151	1399	1647	1895	2143	2391	2321	2639
3	99		465	147	465	651	899	1147	1395	1643	1891	2139	2387	2325	2635
4	95		469	143	469	647	895	1143	1391	1639	1887	2135	2383	2329	2631
5	91		473	139	473	643	891	1139	1387	1635	1883	2131	2379	2333	2627
6	87		477	135	477	639	887	1135	1383	1631	1879	2127	2375	2337	2623
PILOT	84.08		479.92	132.08	479.92	636.08	884.08	1132.08	1380.08	1628.08	1876.08	2124.08	2372.08	2339.92	2620.08
7	83		481	131	481	635	883	1131	1379	1627	1875	2123	2371	2341	2619
8	79		485	127	485	631	879	1127	1375	1623	1871	2119	2367	2345	2615
9	75		489	123	489	627	875	1123	1371	1619	1867	2115	2363	2349	2611
10	71		493	119	493	623	871	1119	1367	1615	1863	2111	2359	2353	2607
11	67		497	115	497	619	867	1115	1363	1611	1859	2107	2355	2357	2603
12	63		501	111	501	615	863	1111	1359	1607	1855	2103	2351	2361	2599
1	107	5	505	107	505	611	859	1107	1355	1603	1851	2099	2347	2365	2595
PILOT	104.08		507.92	104.08	507.92	608.08	856.08	1104.08	1352.08	1600.08	1848.08	2096.08	2344.08	2367.92	2592.08
2	103		509	103	509	607	855	1103	1351	1599	1847	2095	2343	2369	2591
3	99		513	99	513	603	851	1099	1347	1595	1843	2091	2339	2373	2587
4	95		517	95	517	599	847	1095	1343	1591	1839	2087	2335	2377	2583
5	91		521	91	521	595	843	1091	1339	1587	1835	2083	2331	2381	2579
6	87		525	87	525	591	839	1087	1335	1583	1831	2079	2327	2385	2575
PILOT	84.08		527.92	84.08	527.92	588.08	836.08	1084.08	1332.08	1580.08	1828.08	2076.08	2324.08	2387.92	2572.08
7	83		529	83	529	587	835	1083	1331	1579	1827	2075	2323	2389	2571
8	79		533	79	533	583	831	1079	1327	1575	1823	2071	2319	2393	2567
9	75		537	75	537	579	827	1075	1323	1571	1819	2067	2315	2397	2563
10	71		541	71	541	575	823	1071	1319	1567	1815	2063	2311	2401	2559
11	67		545	67	545	571	819	1067	1315	1563	1811	2059	2307	2405	2555
12	63		549	63	549	567	815	1063	1311	1559	1807	2055	2303	2409	2551

FDM Frequency Chart used for making level measurements

Accurate measurement of the many pilot tones required a selective voltmeter or Selective Level Measuring Set (SLMS). This is in effect a tuneable narrow-band filter which selects the pilot to be measured and rejects all the nearby signals and telephone traffic. To do this effectively, a filter of around 30 Hz bandwidth is needed with a tuning range from 60 kHz to tens of Megahertz. The output of the filter is measured by a voltmeter to indicate the pilot level. Tuning over such a wide range requires a super-heterodyne receiver, similar in concept to an old-fashioned radio set.

The test sets used up to the 1970s were manually tuned. The operator would tune the test set close to the required frequency and then make fine adjustments with the large tuning knob to get a peak reading on the moving-coil meter. The gearing and "feel" of the tuning knob was important for ease of use.

The frequency calibration of the test set still needed to be quite accurate otherwise the operator might lock on to the wrong signal. However, the final manual tuning for peak response was essential to eliminate any drifts in the pilot signal itself, in the measurement filter and in the frequency calibration of the test set. Each measurement could take a minute or two to complete.

A high-capacity FDM system might have 300 or more pilot tones, all of which had to be measured to 0.1 dB accuracy, requiring several hours of work. Each pilot has a specific frequency in the FDM plan, so the operator first looked this up on the frequency chart or table and then tuned the SLMS, read the level and noted it down. It was a straightforward procedure but laborious and prone to error. An example of a frequency chart for a 600 channel system is shown on the previous page. Searching complete systems for high-level telephone channels was impractical with a manual test set, so this was only done to diagnose problems on a specific channel or group of channels.

In the era of manual test sets, the German companies Wandel und Goltermann (WuG) and Siemens had fine reputations for making reliable and accurate instruments, such as the WuG SPM-6 18.6 MHz Selektiver Pegelmesser, a favourite with telephone companies, shown here. HP introduced a product in the late 1960s to compete in this market, the 312A Wave Analyzer; however it had small market penetration compared to the German rivals, possibly because it used a rather unconventional homodyne receiver design which produced a notch at the centre of the selective filter. It was also less accurate.

South Queensferry's product would take a new and innovative approach.

A "New-Age" SLMS - The 3745A/B (Unit Preserved NMS T.2010.78)

It is interesting to speculate where the novel ideas for the new SLMS came from. Bob Coackley, a manager in R&D, is often credited with the concept. Bob had worked in Post Office Telephones in earlier years, so was probably aware of the measurement needs.

Around 1971, a proposal for the new instrument was drafted. It was a remarkably innovative idea since it relied on some new technologies, verging on the futuristic at the time. The new SLMS proposal had three significant differentiators:

- It would dispense with the need for FDM frequency tables by storing or calculating all the tuning frequencies, so measurements, or sequences of measurements, could be defined simply with the descriptors of Group, Supergroup, Mastergroup number and so on.
- The tuning of the instrument would use a digital frequency synthesizer locked to a very stable reference oscillator so there would be no need for fine tuning by the operator. The instrument would tune immediately to exactly the right frequency.
- The measurement receiver would be fully auto-ranging and self-calibrating with digital level detection so measurement results could be analyzed and displayed directly by the on-board processor. The measurement filters would be more accurate and stable than earlier instruments so that, in conjunction with the synthesizer, level accuracy was guaranteed without manual tuning.

Design work started in early 1972. The first objective was to build a marketing prototype to test these three concepts and evaluate the ideas with key customers.

The first challenge was to calculate the FDM tuning frequencies from the FDM Channel, Group, Supergroup, Mastergroup numbers. There were thousands of specific frequencies that needed to be defined to 10 Hz accuracy. In the 21st Century, we would simply store these in a gigantic look-up table as semiconductor memory is so cheap today. But in early 1972, there were hardly any memory chips and they were very expensive and stored relatively little information. If there was to be any chance of implementing the idea in a stand-alone instrument, the frequencies would need to be calculated in real time using relatively few stored values. In effect, the algorithms would have to be developed that would mimic the real FDM multiplex system using the hierarchy of multiplex levels and frequency translations.

Related to this, the software designers led by David Dack had to develop a user interface that would be intuitive and easy for the telecom engineer to use. Keyboards on instruments were a novel idea, and it was only three years since HP had brought out the first desk-top calculator, the 9100A. The instrument keyboard needed to be more specific and designed round the measurement functions, while including a standard numeric keypad for data entries. Considerable effort went into defining the necessary key functions and the best sequences to initiate a specific measurement. In particular, the FDM number entries needed to be very explicit so the user could see the FDM descriptor and the calculated frequency.

All this early work was done using the HP 2100 mini-computer, a product that was manufactured at South Queensferry at the time. The mechanical designer, Harry Elder, came up with a neat portable keyboard with LED numeric displays[6] along the top for displaying frequency and level, and a vertical group of LED numbers and associated keys for defining the FDM parameters. There was a number keypad and other measurement specific keys. The keyboard interfaced with the 2100 which in turn drove the LED displays. This allowed the software designers to start work on the FDM frequency algorithms.

The second innovative feature of the new SLMS would be the accurate frequency synthesizer. This provided the tuneable local oscillator for the first stage of mixing in the super-heterodyne receiver. Fortunately, HP had just introduced a new synthesised signal generator, the HP 8660A/B, at the end of 1971. It was the first of its kind, being a digital frequency synthesiser but with the spectral purity of a conventional signal generator. This was important, since a high-performance measurement receiver capable of selecting a specific signal while rejecting all the surrounding unwanted signals, needs a local oscillator with low phase-noise sidebands. The 8660 did this and was also fully programmable through a digital interface. It was simply a matter of writing a driver for the 2100 computer so that it could programme the 8660 with the calculated tuning frequencies.

The third element of the new instrument was the auto-ranging measurement receiver. Here, there was no short cut – the unit had to be designed from scratch in the South Queensferry R&D laboratory. Initially, this was down to the project leader, Reid Urquhart, and myself, Hugh Walker. It was very much uncharted territory, and it is clear from my old lab books from 1972 that we were floundering around somewhat trying to work out how to design amplifiers, filters and balanced mixers that would meet the very demanding analogue specifications of the SLMS. The receiver needed to combine low noise with very low intermodulation distortion while providing measurement accuracy of approximately 0.1 dB over a wide frequency and level range!

As mentioned earlier, the objective was also to incorporate more precise and stable selective filters in the new instrument. These needed to have a very flat response in the passband to guarantee accuracy without manual tuning, while providing high rejection of nearby unwanted signals. This was particularly true of the filter required for selecting a particular FDM telephone channel of 3.1 kHz bandwidth. Up until then, a typical SLMS (such as the WuG SPM-6 mentioned earlier) used a narrower filter (a 1.74 kHz filter was common) in order to achieve the necessary out-of-band rejection for adjacent channels and pilot tones. The new SLMS would have a “true” channel filter with a flat response over 3.1 kHz and very steep sides to reject the unwanted signals.

It was a steep learning curve, but gradually a prototype receiver evolved during 1972. It didn’t have the performance needed for the final unit but the team were gaining experience. In summer 1972, two new graduate engineers, Boyd Williamson and Donald Stewart, joined the design team and work progressed more quickly. By the spring of 1973, the marketing prototype was nearing completion with the first development of the FDM plan algorithms written on the HP 2100 computer. The prototype was a trolley carrying the computer, the 8660 synthesizer and on top the new receiver box and the control keyboard. There was a lot of pressure on the whole team to complete the prototype to demonstrate to some corporate visitors, and several long weekends were spent at the factory striving to complete it. One of the team, brought along a camera, and we have a fascinating set of

[6] Light Emitting Diode numeric displays were a very new invention from HP, as instruments up to that time used “nixie” discharge tubes or moving-coil meters.

photos of the development work progressing. During 1973, it was demonstrated to various customers, in particular some British Post Office engineers who took a close interest in the project and gave many useful inputs.

Building the Lab Prototype

Around this time, work started on a proper lab prototype of what would be the 3745 SLMS. Now a whole lot of new challenges emerged, condensing the hardware and software into a single box.

By the summer of 1973, the final block diagram of the measurement receiver had been agreed. It would have a measurement tuning range up to 25 MHz, and the receiver would have six intermediate frequency stages, taking the measured signal to a low enough frequency to build stable and highly-selective filters. The active 30 Hz pilot filter was centred on 487 Hz, while the 3.1 kHz channel filter was an 11^{th} order elliptic filter centred on 5.55 kHz. It was the first time an SLMS had had a proper channel filter exactly matching the FDM system. Another first for the 3745 SLMS was the Group Filter, a 48 kHz bandwidth filter that would measure the total power in a 12-channel Group. This allowed very fast scanning of the FDM spectrum looking for high-level Groups which might cause overloads and degraded performance.

By then, the design team had built up a lot of knowledge about how to design low-noise and low-distortion amplifiers and balanced mixers. The input mixer which converts the input signal to the first intermediate frequency is particularly critical to the receiver performance. The 3745 used a diode ring mixer driven by a powerful square-wave local oscillator to reduce distortion[7]. The square wave drive, with a rise and fall time of around 1 nanosecond, depended on some proprietary differential-pair transistors from HP, which were also used by the South Queensferry designers of the 3760A pattern generator for the high-speed output stages.

The 3745 receiver had some further innovations. As with all similar test sets, there was the need for range setting. In a traditional SLMS this was typically a rotary switch which allowed the operator to change the receiver gain in 10 dB and/or 1 dB steps to bring the measured signal into the range of the moving coil meter for reading the level. The measured level was the combination of range setting and meter reading. Since the new SLMS was to be fully automatic, making up to three measurements a second, the range setting needed to be done with solid-state switches for speed and reliability, while still maintaining the high accuracy of the mechanical switch.

A further innovation was that the final signal detectors in the SLMS would be "true RMS", that is they would measure the true power of the signal using a thermal detector, in contrast to the diode detectors used on earlier instruments. HP had made RMS voltmeters for many years and in the early 1970s developed some new devices called thermopiles for highly accurate power measurements. The 3745 used this new device and was the first SLMS to make accurate power measurements of any type of signal including tones, noise and wideband FDM traffic signals.

---

[7] "Sources of Intermodulation in Diode-Ring Mixers" by H P Walker, *The Radio and Electronic Engineer*, May 1976, Vol. 46, pp 247 – 255.

Building the first prototype and the completed system in May 1973

The output of the thermopile circuit was a DC voltage directly equivalent to the signal being measured. This linear analogue signal had to be converted to a logarithmic digital bit-pattern to allow the processor to calculate power in decibels. However, multiplication and log conversion were out of the question for microprocessors of this era.

To solve this problem, David Dack and David Arnold devised an ingenious logarithmic analogue to digital converter (ADC). By sending the measured value to the processor in dB, it was then only a matter of adding or subtracting the range settings, reference levels or calibration coefficients in the processor to display the final result.

The ingenious circuit used a voltage to pulse rate converter which generated a burst of 1 MHz pulses; the number of pulses being determined by the received DC level – the higher the voltage the greater the number of pulses. A clever arrangement of up and down counters converted the input pulses into a logarithmic output BCD (Binary Coded Decimal) word sent to the processor. As it was a completely digital implementation, there was no error in the conversion apart from the very small quantizing error. The digital counters could be set to accumulate more or fewer pulses to provide different averaging times or smoothing on the measured result.

The next task was to replace the HP 8660B synthesizer with a custom designed module for the SLMS. By the early 1970s, a lot of work was being done in HP on synthesizer design. In 1973, HP's Loveland Division in Colorado had introduced a new 13 MHz synthesizer, the 3330B, which had fast switching and excellent spectral purity. It looked like it might be a good model for the synthesizer required in the SLMS, so one of the engineers from South Queensferry, John Coster, transferred to Loveland to develop the initial circuitry with the help of the local R&D team. The new synthesizer had a frequency step of 10 Hz, allowing the SLMS to be tuned accurately to 10 Hz resolution anywhere in the frequency range up to 25 MHz. The biggest challenge for John was to minimise the phase noise and spurious signals present in the final output of the synthesizer, as this would degrade the selectivity of the SLMS receiver.

The final step in developing the integrated box was to replace the functions of the HP 2100 computer with an internal microprocessor and memory. These were the earliest days of single-chip microprocessors. Intel's first entry in 1971 was the 4-bit 4004, followed by the 8008 8-bit device introduced in April 1972. The processor team, headed by David Dack, selected the 8008 for the 3745, making the SLMS one of the earliest microprocessor implementations, and probably the first within HP. Apparently it wasn't the initial solution considered, as David Dack recalled:

"*When Bob Coackley dreamt up the SLMS, the idea was we should use the HP35 calculator chip-set as a microprocessor – I think there was some concern that HP's first pocket calculator might not catch on and other uses needed to be found for the chip. I was dispatched to Palo Alto to attend a two-day intensive course on the chip set. I was able to reprogram a Model 35 to calculate frequencies from FDM numbers, but it soon became apparent that it was indeed just a calculator, so thank goodness the 8008 arrived just in time! Kevin Bradford and I went on a training course at Intel. On our return, I prepared a presentation for everyone about the change of plan and was worried about how it would be received. I need not have worried. There was a brief discussion and then Bob Coackley said, 'OK, do it.' That's the way management decisions should be made!*"

The SLMS was one of the first products to claim it had "*Intel Inside*".

The Intel 8008 could be clocked at around 0.5 MHz and could address up 16 kilobytes of memory. Typical instruction cycles took between 20 and 40 microseconds. For its time, the Intel 8008 was an advanced device capable of executing instructions, performing arithmetic, addressing memory and managing data transfer. However its 3500 transistors were contained in an 18-pin package, so the 8-bit data bus had to be used for all transactions including memory addressing. This meant the design team had to surround the chip with various latches and registers to hold addressing and data interfacing to the various parts of the SLMS.

In the 21st century, it is hard to appreciate how limited microprocessors and memory were in the early 1970s. The Assembly Language programming had to be very efficient to minimise the memory requirements and processor instruction cycles. As David Dack recalled,

"*Fortunately Dave Warren joined us in 1973 so we had a real computer scientist on board. This was essential as he wrote an Assembler Program for us as well as joining the hardware team. Intel did not have an Assembler at that time, so it might have been the first one ever.*"

The design team devised various ways of keeping processor time down. One strategy was "pipelining", whereby the processor would be working on the next task while the receiver was processing a measurement (time was needed for auto-ranging and settling of the filters and level detector). An example was the calculation of tuning frequencies from the FDM plan information. When making a multipoint scan of an FDM system, the processor would calculate the next frequency point while the current measurement was in progress, so once completed, the next frequency would be ready to go to the synthesiser without any delay. Another was the use of Direct Memory Access (DMA) which allowed the LED display logic to access the Random Access Memory (RAM) directly to obtain data, so processor time was not consumed servicing the display algorithms.

The marketing group collected all the most important FDM plans from around the world, and one of the software team, Ralph Hodgson, managed to squeeze all of them into very compact list-based data structures ("tables") which, together with a highly efficient algorithm, could be stored in the limited memory available. The final processor implementation used 10 kilobytes of ROM (five 16-kilobit ROMs each providing 2 kilobytes of memory) and five 1-kilobit RAM chips (640 bytes of RAM). During development, the memory card was simulated by the HP 2100 computer. For the first lab prototype, the code was programmed into EPROMs (Erasable Programmable ROMS) and once ready for production, finally into the five mask-programmed ROMs.

As the new SLMS was fully automatic, it was decided early in the project that it would incorporate the new HP-IB[8] interface so the instrument could control peripherals and also be controlled as part of a measurement system. The HP-IB first emerged in 1972, and by 1974 was a useable standard defining the physical interface (16-bit parallel) and the handshaking protocols. The early work on this was led by Don Loughry in the USA and two engineers at HP's Loveland Division. The 3745 SLMS was a very early

[8] Hewlett-Packard Interface Bus, later IEEE-488 and GPIB

implementation, and while in later years custom ICs were available for HP-IB, in 1974 the processor team implemented some of the interface with standard logic ICs and the rest with firmware in the processor, consulting with the pioneers in the USA. The HP-IB software was undertaken by Robin Myles who joined the processor team in 1973.

All this new hardware and software came together into the single instrument in 1974. It was a large box, 10.5 inches high and 20 inches deep, weighing 85 lbs. Most of this space was taken up by the synthesizer in the lower deck and the receiver in the upper deck, both of which required a considerable amount of metal screening which added to the weight. The processor, memory and most of the interfacing logic was installed in a compact housing behind the instrument keyboard, the whole of which hinged out for access and more ergonomic use of the keyboard.

The team quickly discovered that the processor system and the strobed LED displays didn't make good neighbours for a low noise synthesizer and a very sensitive measurement receiver! There was a lot of interference. Eventually the problems were solved by more filtering and screening and improved earth returns. Again it was a new challenge resulting from microprocessor control in the instrument. On the other hand, the processor enabled a range of self-test routines that allowed the receiver and synthesizer to be tested and adjusted.

By the summer of 1974, the lab prototypes were being completed, as I recorded in a personal letter dated 6th June:

"We are just completing our second prototype which will be going to the USA at the beginning of July. The next prototype, which we hope to complete by the end of June, will be taken by David Dack and myself to London during July for a demonstration to the BPO. They have ordered 10 units and we have an order for 7 from Venezuela so that is a total of 17 units, approximately $200k already. David Dack is the project leader for the computer/control side of the instrument – in this respect one the most advanced HP has yet made."

A Convincing Demonstration

We were invited to demonstrate the new product to the British Post Office at their main international exchange at Wood Street in the City of London in mid-July 1974. David Dack, Gordon Reid (from marketing) and myself, spent a week at the exchange demonstrating the prototype SLMS to senior Post Office engineers and managers, using live traffic monitored through protected test points. It was a good opportunity as the exchange was very busy most of the day.

The SLMS could scan the systems fast. In minutes, it uncovered a few pilots that were out of specification and several very high channel levels. There was some initial embarrassment, but the station staff quickly corrected the faults, so as the week went on it became harder to find limit violations during the demonstrations! We left the instrument in the exchange overnight, so it may well have been used to clean up the systems out of hours. When the SLMS found an out of limits signal during a scan, it could halt or could print out the violation as the scan proceeded, using an external printer attached via the HP-IB interface.

It was a convincing demonstration of the power of the new SLMS compared to the old manual test sets. It was possibly a victim of its own success too, as it must have been obvious to the Post Office staff that a few units of the new 3745 could be moved around different exchanges quickly identifying degraded performance, whereas every exchange needed one or more manual test sets.

The new 3745 SLMS with its stored FDM plans and scanning routines could do in minutes what had previously taken hours. It could scan all 275 pilots on a 2700 channel system in 2 to 3 minutes, and all the channels in under 15 minutes. Using its novel 48 kHz group filter, it could scan the whole system for traffic overloads in a minute. In addition to controlling a printer (the HP 5150A) through its HP-IB port, the SLMS could also plot the scanned measurements on a CRT display, so the operator could quickly identify problem signals. This used an optional memory/driver board which slotted in the back of the instrument. Designed by Kevin Bradford, it was a significant piece of electronics given the technology of the day.

During 1975, the new instrument transferred to production. In some ways, the final specification test of the instrument was a simpler task since the 3745 was completely programmable through its HP-IB interface. It was one of the first instruments at the factory to have automated final test, run from one of HP's new desktop calculators, the 9830. The main challenge was establishing the necessary calibration accuracy for the level measurement. Given the very tight advertised specification, the level source used for testing needed to be accurate to about 0.05db over the full frequency range up to 25 MHz and for levels down to -80 dBm.

To verify the accuracy over the level range, a programmable standard attenuator was used. This was developed at the factory and calibrated using a primary standard in the Division's standards lab. South Queensferry worked with the National Physical Laboratory on a special standard attenuator, and the Division was one of the few locations in the UK with this capability.

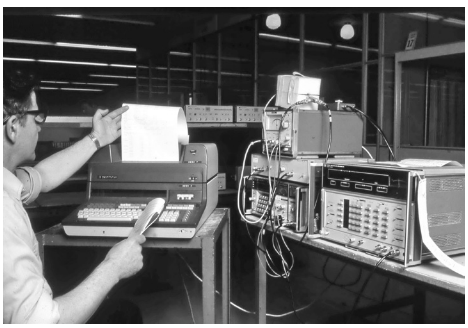

For the absolute level accuracy and frequency response (called "flatness"), a very accurate power meter called a Milliwatt Test Set was needed. HP didn't make one, but our competitor Wandel und Goltermann did. The WuG EPM-1 was something of an industry standard. Division management was initially somewhat reluctant to buy from a competitor, but it proved the right decision and several were bought over the years and used in R&D and manufacturing test.

Not only did this novel automatic instrument require some fresh thinking in production test, it also presented a challenge for servicing and post-sales support. HP instruments were all supplied with detailed service manuals so that the products could be repaired to component level in HP service offices round the world and also by knowledgeable customers. At that time, the new digital electronics, and particularly the Intel microprocessor in the 3745, would be unfamiliar to service engineers, and probably quite daunting. The microprocessor did provide some advantages with various self-test and alignment routines for the measurement hardware, however if the instrument's "brain" was faulty, life got complicated! Early on, the Queensferry product support engineers realised this would be a problem and devised a special diagnostic kit (15585A) to assist the service engineers troubleshoot the processor system. Garry Irvine, one of the 3745 Support Engineers, describes this system in Appendix 7.

First shipments took place in February 1976 with a delivery to the British Post Office, part of an initial order for 10 units. Quite a lot of negotiation took place regarding the price which was two or three times what the BPO would normally pay for an SLMS. The 3745 had a hefty US list price of $21k at launch, although this was reduced to $17k by 1978. Two versions of the 3745 were sold. The A model was for the European/International market with CCITT FDM plans (shown here with display and printer), while the B model was for the North American market and contained the Bell System FDM plans and WECO signal connectors.

It was not a big seller. The combined “A” and “B” models sold around 8/month from 1976 until 1981 when it was replaced by the 3746A described later. A total of 374 “A” models and 144 “B” models were shipped with total revenue of around $10M.

Given the enormous advantages over the manual SLMS, the Division had expected sales to be higher. Several factors might have influenced the outcome. The 3745 was very expensive, so telecom operators might have wondered if they could justify the cost on productivity gains. In the 1970s, most telephone systems were state-owned and highly “unionised”. The 3745 would automate procedures and could result in lost jobs, so it was another issue to be considered. Finally, by the late 1970s digitization of the telephone network was underway, which would eventually sweep away FDM transmission. The very technology that had made the 3745 possible, digital electronics, was bringing in a revolution. While some operators just stuck with what they knew, others decided to go for complete automation in order to reduce costs with the legacy analogue transmission, hence the growth in FDM surveillance systems discussed later.

The design of the 3745A/B is described in more detail in the *HP Journal* of January 1976[9], which features the instrument on the front cover. The unit preserved in the National Museum of Scotland (NMS T.2010.78) is a 3745B manufactured in early 1980. The story of its repatriation and restoration for the Museum is told in a postscript at the end of this chapter.

---

[9] HP Journal, January 1976 pp 2 – 18. http://www.hparchive.com/Journals/HPJ-1976-01.pdf

The 3747A/B 90 MHz SLMS

In late 1975, I was on a business trip to the USA with the then R&D Manager, Bob Coackley. We visited Bell Labs at Holmdel in New Jersey where the engineers were keen to tell us about their new L5 and L5E 60 MHz FDM cable systems. Afterwards, as the local HP field engineer, Joe Arcidiacano, drove us back to the hotel, I remember Bob turning round to me and saying, "*So, Hugh, can we make a 100 MHz SLMS for these guys?*" "*Well, I've been thinking how we could make a 70 or 80 MHz SLMS.*" Over the years, it was sometimes the case that a visit to a customer would spark off the beginnings of a new idea, which is why HP was keen that its R&D engineers visited customers.

The idea in this case was quite a simple one. We would roughly double the 3745's first intermediate frequency (from 61 MHz to 116 MHz) and fit a frequency doubler on the output of the frequency synthesized local oscillator. The receiver arrangement wasn't quite as well optimised as the 25 MHz SLMS and the filtering would need to be tighter, but it would probably work well enough. And that was the genesis of the 3747 SLMS[10].

The final design went to 90 MHz with a reduced specification, while the performance to 60 MHz was nearly as good as the original SLMS. David Haworth modified the synthesizer to have a 5 Hz step size (converting to 10 Hz after doubling), extended the tuning range and built a frequency doubler. Gregan Crawford augmented the stored FDM plans to include the various 60 MHz systems in North America and Europe. The processor board was redesigned to include space for an additional optional EPROM memory to house customised non-standard FDM plans, in addition to the mask-programmed ROMs used for the rest of the firmware. Shipping instruments with EPROMs was a new innovation, later to become standard; nobody was really sure that these erasable programmable read-only memories would retain their programmed images so there were certainly a few crossed fingers! This processor board was then used in both the 3747 and the 3745.

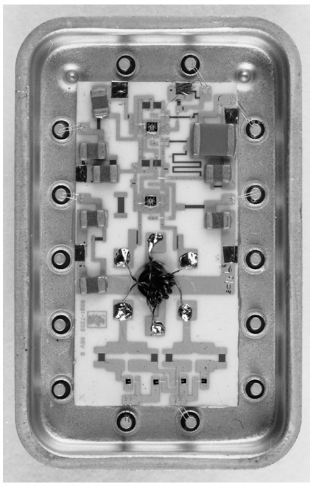

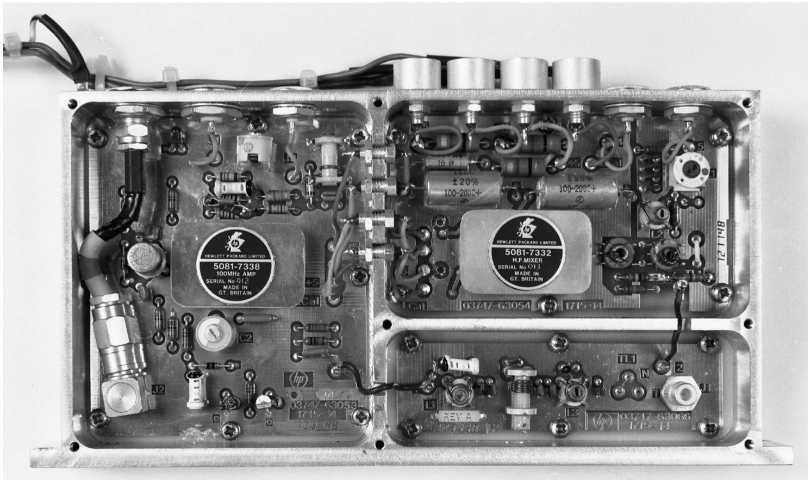

I redesigned the receiver front-end, taking advantage of the new thin-film circuit technology developed at South Queensferry. It replaced printed circuit boards with an alumina substrate and used chip components, which roughly doubled the bandwidth and performance. The photos above show the input balanced mixer hybrid and the new front-end module. The thin-film technology had already been used in the new Microwave Link

[10] The original plan had been to go to 160 MHz, but the Bell Labs visit gave the impetus for a simpler, quicker and cheaper solution.

Analyzer and in the 3762/63 Pattern Generator and Error Detector. In effect, the new 3747 was the original 3745 front end transferred to thin film and with tighter filtering. In other respects it was very similar to the earlier instrument.

Shown here with the companion 3335A Synthesizer, first shipments of the 3747A/B took place in February 1978. With a US list price of around $21k, it looked attractive as the competitive instruments from Wandel und Goltermann and Siemens were higher priced or of lower specification[11]. Sales volume was relatively low, around 40/year up until 1981, mainly due to the relatively small number of 60 MHz systems installed. The 3747 was also popular with the final generation of high-capacity analogue submarine cable installations in the early 1980s. Gregan Crawford implemented a number of special FDM plans for these customers and the instrument had an optional 2.5 kHz channel filter for testing the 3 kHz channel spacing used in these submarine systems. The instrument was finally discontinued in 1985, a total of 181 units being shipped with revenue of $5M. The price was raised significantly over the years, reaching $33k in 1985.

The Access Switches (Unit Preserved NMS T.2013.63)

With the built-in speed and automation of the 3745 and 3747, it was logical to think of them being used with programmable access switching so that a large number of test points in the transmission centre could be monitored sequentially. The switches could be positioned next to the transmission equipment test points and the signals routed back to the SLMS which could be installed at a central control desk or form part of a remote monitoring system.

The first switch to be developed was the 3754A, a 10-to-1 switch operating to 25 MHz to partner the 3745 SLMS. It used a wideband shunt-feedback amplifier with a novel arrangement of relay switching at the amplifier's virtual earth, which reduced the effect of stray capacitance. This had several advantages in terms of impedance matching and isolation between channels, and was the subject of a UK Patent[12]. The amplifier provided a small amount of gain that could be used optionally to offset losses in cables and equalisers.

Three hierarchical levels of switching would allow any one of up to 1000 separate inputs to be routed to the SLMS, using a pyramid of 3754 switches. The challenge then was how to programme these switches, as having an HP-IB control interface on every switch was impractical and far too expensive.

David Dack came up with an ingenious control system that mimicked the way old electro-mechanical telephone exchanges worked. Low-speed pulses, similar to the old telephone dial pulses, were applied to the output port of the first switch in the tree. The 3754A detected these and used them to select one of its 10 inputs. This connected to a switch in the next layer, and once the connection had been made, the next set of pulses was

[11] The WuG PM-7, a multi-box solution costing over $30k was the most often seen.

[12] British Patent 1 482 290 "Improved Switching Means", Filed 5th August 1974, published 10th August 1977.

automatically routed to the next switch, and so on. It was a delightfully simple solution, since once the signal coax cables had all been connected up, a control path also existed. The whole thing worked, because the switch's wideband amplifier was AC coupled and didn't respond to the low speed dialling pulses[13].

Obviously there needed to be a source of dialling pulses, and this was done by connecting the output of the final switch through a unit called the 3755A Switch Controller. It had an HP-IB interface and converted control commands into dial pulses which it transmitted to the first switch.

The 3754A 25 MHz Access Switch was introduced in April 1977, and cost around $2k, while the 3755A Controller cost $1.5k. The 3754A was in production until 1991 and a total of 2100 units were shipped. The 3754A and 3755A were described in an article in the *HP Journal* issue of August 1978[14]. (A 3754A is preserved in the National Collection, NMS T.2013.63)

One follow-on project after the 3745 was a lower-cost SLMS working up to 8.5 MHz suitable for 1800 channel FDM systems. Designated the 3748A, work continued on it in the late 1970s, although there were some questions about its viability. Growth was slowing in the FDM market, although it continued to be a large one in dollar volume, and competitors such as W&G introduced more new products. There was also a charter conflict with another division who were building a lower cost SLMS (discussed later). For whatever reason, development on the 3748A was discontinued, but part of the package was a low-cost access switch also going to 8.5 MHz, and this continued into production as part of the access switch family. This was the 3757A and was similar to the patented concept of the 3754A, except that solid-state switching replaced relay switches for channel selection, albeit with lower dynamic range. It used the same "dialling" process for channel selection and was compatible with the earlier switch so they could be mixed in a system. With first shipments in December 1978, the 3757A sold for around $700, increasing in later years. During its 12 year production life, over 1000 units were built.

The final product was the 3756A 90 MHz switch, introduced to partner the 3747 90 MHz SLMS. It was a straightforward DC-coupled relay switch so could operate as a 10-to-1 access switch or a 1-to-10 distribution switch. Like the other access switches, it was designed for 75 ohm applications in telecom. HP made other coaxial switches but mainly of 50 ohm impedance for RF and microwave applications. So, the 3756A was used in other non-SLMS systems over the years where 75 ohm impedance was required. It was also controlled using the same principle as the 3754A, however only over the 2-wire alternative, not via the signal path. The 3756A was introduced in July 1978, and sold for around $2.5k initially. Around 900 units were shipped until discontinued with the other access switches in 1991. In total around 4000 access switches were sold, bringing in about $10M.

[13] The system could be set for low speed (20 pulses per second) or high speed 600 pps. There was also the option to route the control pulses through twisted pairs so that they could be isolated from the signal path, if desired.

[14] "Remotely-Controlled RF Switch for Multipoint Tests in Communications Systems" by Kevin Bradford, HP Journal August 1978 pp 20 – 22 http://www.hparchive.com/Journals/HPJ-1978-08.pdf

A Charter Conflict and a New SLMS (or Two) – The 3746A SLMS

A key philosophy of Bill Hewlett and Dave Packard was to encourage innovation and a feeling of ownership among employees. They wanted to create a small company ethos within a worldwide corporation. During the great era of expansion from the 1950s to the 1980s, growth was achieved by creating a series of autonomous product-focussed divisions, with never more than a few hundred employees each. The management in these divisions was responsible for running a business, where growth was financed from the locally generated profits. It was a self-contained business (product development, marketing, production), a profit centre, and there was a lot of freedom to develop product strategy. The main restriction on management was that they had to stay within their product charter, such as frequency counters, oscilloscopes, signal generators, or whatever.

This delineation worked well in the early days, but as the business grew, the market became more congested and overlaps occurred. Furthermore as technology advanced, particularly in the 1970s, it became economic to combine measurement capabilities, and indeed customers expected it. So charter conflicts became more of an issue, and it usually meant Hewlett and Packard and the corporate team getting involved.

For many years, South Queensferry had very few charter conflicts. Nobody else was that interested in the telecom test business, indeed many of the other HP divisions had no idea what it was all about. However, one division, Loveland Division in Colorado, did have an interest. From the 1960s, they had responsibility for Wave Analyzers which is another name for Selective Level Measuring Sets. In 1967, they introduced the 312A Wave Analyzer covering the frequency range up to 18 MHz and with filters suitable for telecom measurements. Presumably they hoped to compete with companies like Wandel & Goltermann and Siemens. In 1970, they introduced the 3556A Psophometer, a fancy name for a tester that measured noise and signal levels on telephone circuits. Finlay Mackenzie recalled that in the late 1960s, Loveland's then Engineering Manager, Bill Parzybok, campaigned to have South Queensferry's telecom design team, headed by Peter Carmichael, transferred to Colorado, but Bill Hewlett and Dave Packard disagreed and Scotland kept the telecom business.

The two divisions did work together quite closely. As mentioned earlier, the synthesizer in the 3745 SLMS was derived from a Loveland design, and much of the early work on HP-IB was done there too. However, when the new 3745 SLMS was launched in 1976, Loveland responded with a slightly upgraded 312D 18 MHz Wave Analyzer with some new filters and targeted at the telecom transmission market by providing the appropriate connectors and balanced inputs. It sold for $5400, less than a third of the price of 3745.

It was a clear signal that they didn't want to give up this market to South Queensferry. Despite the upgrade and the telecom features, the new 312D remained an old-fashioned manual instrument with various rotary range-changing switches and tuning knobs. Its level accuracy was inferior to the 3745 and the old W&G manual boxes. As mentioned earlier, it used a homodyne[15] circuit for implementing the selective filters and this necessarily

[15] The homodyne circuit used two quadrature (90 degree phase shifted) mixers down-converting the signal to DC where low-pass filters provided the selectivity. It was then mixed up to the IF again to provide the bandpass characteristic. It was necessary to eliminate the DC component as it would disturb the measurement and this led to the notch in the receiver response

meant there was a narrow notch at the band centre. To make an accurate level measurement, the receiver needed to be detuned slightly either manually or using the receiver's automatic frequency control. This was an unfamiliar procedure to users, and slowed down measurements. Loveland Division decided it was time for a new Wave Analyzer or Selective Level Meter (SLM).

In the mid-1970s, Loveland developed a new synthesizer technology called Fractional-N synthesis[16]. The circuit design for this system was challenging, but the block diagram was simpler and more compact. The first product to use this new method was the 3335A 80 MHz Level Generator in 1978 – a first class product, noted for its frequency resolution and level accuracy. It was logical that they would use this synthesizer circuit in a new SLM.

The new product would also be more like a conventional instrument with regular selective filters and a tuning knob and meter, though programmable through HP-IB. The receiver design was simplified by using a crystal filter that eliminated several mixing stages in the receiver. The new product was the 3586 Selective Level Meter[17], covering the frequency range to 32.5 MHz.

There was an obvious conflict of interest between this development and the products from the telecom division at South Queensferry. The conflict became acrimonious at times – R&D engineers and managers can be rather protective and territorial about their inventions! So, who owned the market for Selective Level Measuring Sets? One thing was certain – South Queensferry definitely owned the FDM Plan algorithms.

A compromise was reached whereby Loveland would produce the basic SLMS and Scotland the version with the built-in FDM plans focussed on FDM surveillance, so HP would have all applications covered. A new development started at South Queensferry in parallel with the 3586 development in Colorado. This would be the 3746A SLMS.

A small project team led by Boyd Williamson, one of the young engineers who joined the 3745 project in 1972, took the 3586 design and modified it. Some sections such as the Fractional-N synthesizer were used more or less directly, as well as the simplified receiver structure.

The designers transplanted the high-performance front-end from the 3745 for better dynamic range, and redesigned the back-end measurement hardware with more auto-ranging steps to improve level accuracy. One challenge was to include the 48 kHz group filter that proved a success in the 3745. The 3586 receiver design had insufficient intermediate frequency bandwidth, so a separate receiver path had to be developed for this measurement using a second crystal filter.

The major change was the incorporation of the FDM plans. The knob and meter in the 3586 were discarded, and instead push buttons added for setting FDM plans. The frequency display could be switched between frequency and FDM description. By this time, processor memory had advanced to the point where all the FDM plans could be included as standard. The 3746A also incorporated access switch control so it was no longer necessary to use the 3755A Access Switch Controller, described earlier.

[16] A key inventor of this was Chuck Kingsford-Smith (US patent 3928813A, filed in 1974).
[17] Described in HP Journal, May 1980 http://www.hparchive.com/Journals/HPJ-1980-05.pdf

Overall, the end product was well-differentiated from the 3586[18]. It had pretty much the same functionality and specifications as the 3745 but with slightly faster measurement speed. The designers added some further enhancements. One was a Supergroup power measurement, effectively a software addition of five group power measurements, useful for checking total traffic power. Also, at the request of a US customer, high-speed scanning algorithms were developed to detect high-level signals in seconds rather than minutes. With a view to the FDM systems applications discussed later, the 3746 could store a set of equalization values so cable losses could be corrected automatically in firmware rather than with physical cable equalizers. Like its predecessor, an output to drive a CRT display, the 37461A, was available for graphical presentation of scanned results.

First shipments of the new SLMS took place in September 1980. Of the early orders, 75 units went to the UK, probably British Telecom (Tester 244A) following on from their earlier 3745 orders. When launched, the US list price was around $18k depending on options, a substantial premium over the $11k charged for the 3586 SLM from Loveland. Presumably this affected sales, and by the mid-1980s the price had been reduced to around $13k (US list).

With its newer technology, the 3746A was less than half the size of the earlier 3745 SLMS, being 7.5" high and 19.5" deep. However the large amount of internal metal screening needed to maintain the low noise floor, plus the analogue power supply, resulted in a heavy box weighing 55 lbs (25 kg).

The eventual $13k price tag and smaller size made the 3746A more attractive than the bulkier and more expensive 3745. Typical production volumes were around 12 units/month. Sales were spread fairly evenly around the world, with the highest share being in North America.

In 1987, the Division won a major $6M contract against stiff competition to supply 300 units to AT&T in a very short timescale for their North American transmission network. Success was partly down to the product's capability and fast delivery, but also due to the HP Company being able to provide a complete solution of repair/loaner and on-site calibration services, ensuring high availability in the field. 50 of the 300 units were specified to measure up to 80 MHz, for the L5/L5E cable systems. South Queensferry designed a special additional instrument that worked with the 3746A to give extended frequency coverage. Designated 3746A Opt. H40 Extended Frequency Module, it appeared to be similar in concept to the 3747 90 MHz SLMS of 10 years earlier. Guy Douglas who worked on this module recalled that the development was done in Production Engineering as by then R&D had gone completely "*digital communications*".

The first units were shipped in early 1988, and the contract completed six months later with the shipment of the Opt. H40 80 MHz version. It must have been one of the last

[18] Loveland did have one final trick. They put the FDM plans in an HP85 desktop calculator to control the 3586, the system being marketed as the 3046, coincidentally one digit different to the 3746!

major contracts to supply FDM test equipment, and quite surprising as around the same time AT&T had introduced a 1.7 Gb/s optical fibre transmission system! It did demonstrate that the analogue telecom network was still big business in the 1980s, despite the march of digital. The H40 system is shown here.

In total 1500 units were shipped with revenues of nearly $25M. Around 150 units of the 37461A were produced with a list price of $3.5k. The 3746A was discontinued in 1991 along with all the access switches, bringing to an end the FDM product family. A 3746A has been preserved, and it appears to be one of the loaner units used to support the AT&T contract.

FDM Surveillance Systems

HP entered the computer market in the late 1960s with its first mini-computer, the 2116A. The rationale was that they needed a compact dedicated instrumentation computer for building measurement systems. A few instruments had remote control interfaces, one example being the 8660 synthesizer mentioned earlier in this chapter. In the early 1970s, a systems integration group was set up at South Queensferry to bring together instruments from different divisions under common computer control with customised software for particular applications.

Once the Hewlett-Packard Interface Bus (HP-IB) control had been developed by the mid-1970s, system integration became easier because of the standardised control interface and handshaking protocols. HP-IB compatible products were daisy-chained using HP-IB cables, with transactions to and from a specific instrument being handled through addressing.

The 3745 SLMS was one of the earliest HP-IB enabled automatic instruments, so systemisation was an interesting opportunity that caught the imagination of David Dack, the leader of the 3745A microprocessor development. It was a very new idea, as hardly any other SLMS products on the market at the time were programmable, and those that did have remote control, had a very basic parallel interface requiring a lot of programming. With so much automation already built into the SLMS (FDM frequency calculation, multipoint scanning, limit testing on measured levels), the role of the system computer was more one of storing FDM system information by access point, initiating a measurement scan, and recording the measured levels to form a record of the system performance at a given time.

Sometimes it was desirable to have a complete record of, say, all pilot levels on a system, but normally only those signals out-of-limits would be sent back to the controlling computer. With up to a possible 1000 separate monitor points connectable through the access switching, a colossal amount of measurement data would rapidly accumulate, so an important task for the computer software was to filter, store and present this in a useable form.

The first computer-controlled SLMS monitoring system was developed in 1976 and installed with BC Hydro, in Western Canada, in September of that year. The system comprised a 9830 desktop computer, located at the company's control centre in Burnaby near Vancouver, and a number of 3745s installed at power stations in the north of the province. These were used to monitor the telecommunications links controlling the company's power plants. This system configuration was subsequently used to demonstrate the concept of remote SLMS monitoring to prospective customers worldwide. A key outcome was the interest shown by the North-East Area of AT&T Long Lines in the USA.

Work started on a computer-controlled surveillance system in spring 1977 when David Dack and Dave Warren, both of whom had worked on the 3745 processor, transferred to an HP sales office in Paramus, New Jersey for six months. Here they worked on prototype software in conjunction with AT&T Long Lines staff[19] at White Plains, New York. This system used the HP 1000 mini-computer, and the project was for the AT&T Distributed Maintenance System (DMS) which was the AT&T Long Lines version of the Carrier Transmission Maintenance System (CTMS). The software would be supported by AT&T and the measurement hardware by HP. The plan sounded simple enough, put there were problems writing a service-support contract across HP sales regions which in turn didn't match AT&T's organisation. It was South Queensferry's first experience of the problems of large-scale systems sales in the HP organisation, problems that would surface again with future systems developments.

Two versions of the software evolved in the late 1970s. One was developed to run on HP's new family of desktop calculators/computers, initially the HP 9825 and 9835/45 and later the 9816. This version was intended for smaller systems, typically controlling one SLMS with some access switching, but could control up to six SLMS systems in its final version, albeit sequentially. Although the computing power of these early desktops was limited, because the 3745/47 SLMS did so much of the measurement routine semi-automatically, it was a very workable and cost-effective solution. This smaller system had the product number 37051S and the software package was 37014A, becoming a B/C version in the 1980s to run on HP's developing family of desktop computers.

Larger and more powerful systems needed the mini-computer. This system had the product number 37050S and ran on the HP 1000 A-series mini-computer. The computer-based system could control up to 16 SLMS measurement systems each of which could in theory have up to 1000 access points. The key factor with the computer system was the use of HP's Real Time Executive (RTE) operating system which made the computer multi-tasking. Whereas the desktop version had to control the measurement subsystems sequentially, the RTE software allowed all 16 measuring sets to operate concurrently making a much more powerful and efficient system. The computer and its associated disk storage meant it could also handle and process the large amount of measurement data being generated. The software package for this system was the 37013A, introduced in 1979, and the later version (around 1984) was designated 37016A, which operated with the new 3746A SLMS.

The software could produce a variety of reports depending on the customer need. The example shown on the next page is a status report for the group pilot levels. By accessing

[19] Robin Myles recalled the key guys at AT&T were Andy Litosch and "Buffalo Bill" Foster.

previously stored measurement data, the system showed how limit violations changed over time:

MONITOR MODE PRINTOUT

SQD1::12 PUINT1 11:16 02 NOV 1979 GRPM

NOM LEV = -20.00 UPPER LIM = 2.00 LOWER LIM = 2.00

FDM			FREQ	LEV(DBM)	LEV(DBMO)	
4	1	0	1028.08	-52.70	-19.70	
4	2	0	980.08	-53.50	-20.50	
4	3	0	932.08	-52.50	-19.50	
4	4	0	884.08	-52.70	-19.70	
4	5	0	836.08	-52.90	-19.90	
5	1	0	1276.08	-81.70	-48.70	STILL OUT OF LIMITS
5	2	0	1228.08	-52.90	-19.90	
5	3	0	1180.08	-52.80	-19.80	
5	4	0	1132.08	-53.00	-20.00	
5	5	0	1084.08	-52.90	-19.90	
6	1	0	1524.08	-52.90	-19.90	
6	2	0	1476.08	-52.90	-19.90	NOW WITHIN LIMITS
6	3	0	1428.08	-52.80	-19.80	
6	4	0	1380.08	-52.90	-19.90	
6	5	0	1332.08	-52.50	-19.50	
7	1	0	1772.08	-52.70	-19.70	
7	2	0	1724.08	-52.70	-19.70	
7	3	0	1676.08	-52.70	-19.70	
7	4	0	1628.08	-52.70	-19.70	
7	5	0	1580.08	-52.70	-19.70	
8	1	0	2020.08	-52.60	-19.60	
8	2	0	1972.08	-52.60	-19.60	
8	3	0	1924.08	-90.40	-57.40	NOW OUT OF LIMITS
8	4	0	1876.08	-52.60	-19.60	
8	5	0	1828.08	-52.60	-19.60	

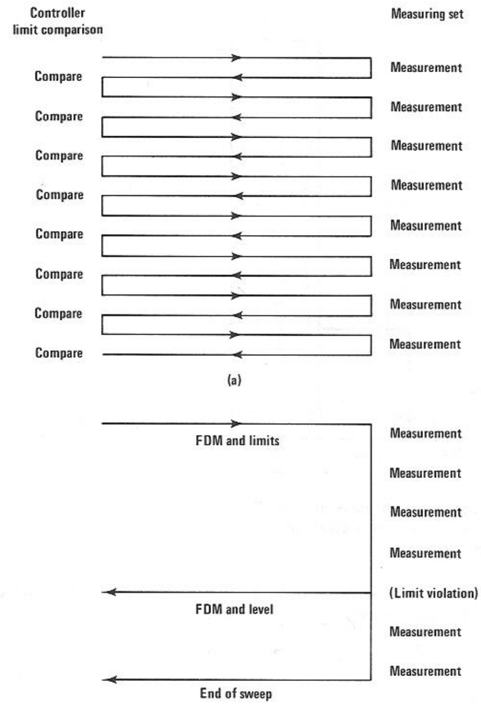

As already mentioned, the built-in measurement routines of the HP SLMS instruments, meant the transactions between the computer and test systems were greatly reduced, as shown here. With a basic SLMS, every measurement point would need to be programmed from the computer and the level measurement result returned to be checked against the limits. With the HP SLMS shown in the lower graphic, very few transactions are required. This had the benefit of faster measurements and each computer could control more measurement systems. There was an even greater advantage if the measurement system was at a remote location and being controlled over a data link. In the late 1970s this would have used a voice data modem operating at 1200 or possibly 2400 bits per second, adding considerably to the measurement time if frequent transactions were required.

In order to control a remote measurement station, the parallel HP-IB interface had to be converted to a serial RS-232 bit stream to drive a data modem. HP had introduced a box to do just that around 1975, the 59403A Common Carrier Interface or CCI (common carrier because it was intended to work over telephone lines). The system designers quickly ran into a problem. The CCI did not incorporate error correction, and bit errors were prevalent on those early modems. It could detect errors but not correct them, so the whole HP-IB interface would hang-up. Not an acceptable situation for a remote monitoring system that had to work reliably 24 hours a day!

Necessity is the mother of invention, so David Dack and the systems team decided to develop an HP-IB Extender that would work reliably. A new product, the 37201A Bus Extender was born. It had error correction, and more attention was given to the handshaking protocols at each end of the link. This successful product led to a whole family of HP-IB Extenders, and over the next 20 years tens of thousands were shipped from South Queensferry, becoming the de-facto standard for HP-IB networking. The remarkable story of these products is told in Chapter 12.

As for the FDM surveillance systems, the sales were rather disappointing with only a few major systems installed round the world, including some smaller systems on submarine cables[20]. Despite this, systems were obviously seen as a future opportunity judging from comments by Finlay Mackenzie, the Divisions new general manager, in the April 1982 issue of the Readout newsletter:

"*In the Christmas Readout I indicated that future trend is towards instruments controlled by central computers. Already we have installed surveillance systems on three European submarine cables and recently won an order for the TAT-6 and TAT-7 Trans-Atlantic cables. The total value to the HP Corporation of these test systems is around $800k.*"

Nevertheless, the revenue from these was totally dwarfed by the income from the HP-IB Extender products that were born out of the FDM surveillance system development!

Although overall system sales were disappointing, one region where the Division did have success was Canada. This was probably due to the efforts of Bill Lauchlan, a South Queensferry test engineer who worked on the Microwave Link Analyzer production line in the early 1970s and later transferred to the USA joining the sales team in Canada in the late 1970s as a systems engineer. He understood the needs of the customer and had in-depth knowledge of HP products and systems software. Working closely with the Robin Myles and Dave Warren in the Division's systems group, they sold the first desktop-computer controlled system in late 1976 and the first HP1000 computer system in 1978.

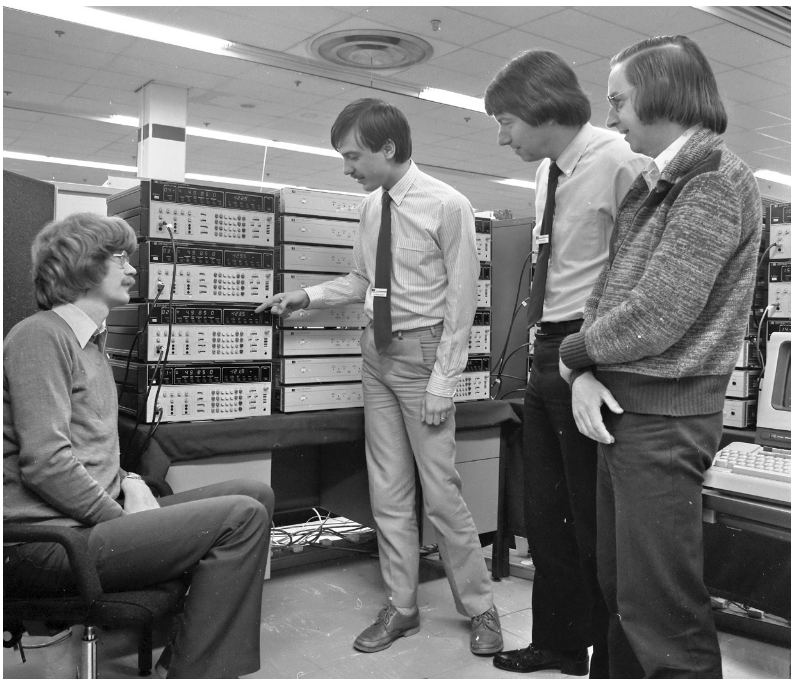

A stack of 3746 SLMSs and 37201 HP-IB Extenders being tested for an FDM Monitoring System
The software design team (L to R): Dick Middleton, Ray Scott, Robin Myles and Vince Butler, early 1980s

[20] The European multi-cable submarine monitoring system was a significant contract commissioned in September 1980, with UK locations at Lowestoft and Land's End and foreign locations in Denmark, Holland and Spain. These were high-capacity cables and the system used the 3747A 90 MHz SLMS. The systems also incorporated the HP 2240A process controller to integrate alarms from the cable itself.

Many more systems followed for both the Canadian domestic and international networks, including satellite and submarine cable links in the following 10 years (see Appendix 8). As Bill Lauchlan later recalled:

"*It was a big learning experience about telephone company politics for both Jack Weldon (Sales Engineer) and myself. Although we didn't get some of the big deals because of politics, the business generated a lot of interest in Palo Alto, and Jack and I were summoned by Finlay Mackenzie and Peter Carmichael (South Queensferry's GM) to meet Hal Edmondson (their corporate boss) and brief him on the whole thing. I think that was important to get Hal on-side for future systems business.*"

There is no doubt that the combination of powerful measurement instruments and the computer software did the job of FDM surveillance superbly well – there was nothing else like it on the market. So why did the business fail to materialise? There were probably two factors that contributed to poor sales.

Larger surveillance systems required a lot of planning and a considerable capital investment by the telecom operator, so an economic justification had to be made based on operational savings over a period of time. By the early 1980s, FDM's days were numbered and many operators were focussing investment on digital facilities. Maybe the appetite for large investments in the analogue network was no longer there.

But perhaps the biggest obstacle was the HP sales company. It dealt directly with customers, but was ill-equipped to deal with the contract negotiation and the time-consuming challenges of large systems sales. Most HP sales engineers were "box salesmen" because what HP sold were stand-alone instruments in the main. Big deals were usually multiple purchases of the same product or bundles of instruments. Even the computer controlled test systems HP supplied were racked and tested in the factory and shipped as turnkey systems.

The FDM surveillance system on the other hand was distributed, had to be integrated into the customer's operational equipment, possibly with customised software, and became critical to their network operations. Who, for example, would install the system and run cabling in the exchange, and who would commission the complete system and hand it over to the customer? Nobody in the conventional HP sales organisation did that kind of work. Inevitably, the Division got involved and the few successful sales happened to be where the local HP sales management shared South Queensferry's vision of the systems potential and was willing to accept a longer sales cycle, as was the case in Canada.

The Legacy

Although the Division didn't find the "Holy Grail" of systems sales with FDM surveillance, the management could see the untapped potential and learned a great deal through the experience. This led in the early 1980s to the successful RATES system described in Chapter 7, and later to the even more successful monitoring systems for Signalling System No.7 which eventually led to the spin-out of a new division in the 1990s.

So, the legacy of the SLMS and FDM Monitoring Systems was not so much in their innovation and sales revenue, but more in what they led to. There is no doubt that without them South Queensferry would never have produced the HP-IB Extenders, which amazingly in the late 1980s were responsible for nearly 30% of the Division's profits (see Chapter 11). Without having seen the potential and understanding the challenges of system sales, it is possible South Queensferry might never have developed the RATES system or the No.7 Monitoring Systems.

Acknowledgements

The following former HP employees contributed to this chapter: Finlay Mackenzie, Boyd Williamson, Gregan Crawford, Robin Myles, David Dack, Garry Irvine, Bill Lauchlan and Guy Douglas.

Bringing a 3745B back to Scotland for the National Museum

In the summer of 2010 as the South Queensferry factory finally closed down, we assembled a small collection of instruments representing the products developed at the site over the 40 years from the late 1960s to its final closure. These were to be donated to the National Museum of Scotland. One important item we lacked was a 3745 SLMS, and we felt it was valuable to acquire one, given its significance in the history of HP at South Queensferry

About this time, several of us noticed a 3745B advertised in an on-line site called Bonanzle (similar to eBay) for just $50. It was located in Regina, Saskatchewan in Canada. The unit did power-up, but had some cosmetic damage requiring restoration and appeared to have electrical faults when tested. We decided to buy it, and Agilent agreed to ship it back to Scotland.

It was shipped first to Toronto, where it was then forwarded back to South Queensferry, arriving around the 22nd July 2010, pretty much the last day of regular operations at the site. The packing for the instrument from Toronto to the UK was inadequate and the instrument had a very rough ride home suffering some additional damage – the worst three days of its 30 year life! Weighing 85 lbs, each time the instrument was dropped it would have given quite a jolt to the electronics inside.

We don't know the full history of the unit, but it was most likely bought and used in Canada. It is also most likely to have been used in an FDM surveillance system. Evidence for this is the presence of rack-mount flanges when bought, the lack of the CRT display driver card (normally preferred for manual use) and finally the controller switch behind the keyboard being set to "External" (for systems use), rather than "SLMS" (the normal setting for stand-alone use).

The instrument is a fairly standard HP 3745B designed for the North American market with Bell FDM Plans and large WECO connectors[21]. Serial number 2010U00222

[21] WECO – Western Electric Company (Bell Manufacturing) connector Type 477B

indicates it was manufactured in the UK in 1980 sometime after Week 10. It was No. 122 of the complete production run of 3745Bs which totalled 144 units. So today this is quite a rare instrument. Printed circuit assemblies in the instrument are dated April/May 1980, confirming the approximate date of manufacture.

I took the instrument home for restoration and repair. The main areas of physical damage were a bent keyboard latch assembly on the right, and a broken filter switch in the top centre of the keyboard. I took out both keyboard latches and repaired, cleaned and lubricated these with silicone grease. At the same time I cleaned the keyboard panel and the keys. The most difficult repair was the filter switch which had been snapped off inside. I re-attached the toggle with Araldite reinforced by a small piece of aluminium. These three photos show the inadequate packaging, the damage, and the restored instrument.



We were very fortunate that the original side-handles were with the instrument, and when the unit was decommissioned, it appears that the handles were reversed behind the rack-mount flanges for storage/transport. This was fortuitous as the handles are easily damaged, and the rack-mount flanges had probably helped to protect the instrument during its rough journey home judging by the damage to them. After restoration the handles were refitted to give the instrument its correct appearance.

The SLMS was not operating properly when received. The keyboard and processor appeared to be working correctly as was the synthesized local oscillator assembly. I cleaned the fan filter and removed an obstruction from the fan blades so that it ran smoothly.

The problems were in the receiver chain. Firstly, an active notch-filter at the back end was oscillating for some reason and preventing the instrument from auto-ranging. Its role seemed peripheral to the instrument operation, so I disabled it. The instrument then calibrated and auto-ranged correctly. The second problem was a high level of intermittent noise in the IF chain. I located this to the Programmable Attenuator assembly, and to a noisy transistor which I replaced.

The SLMS then operated correctly with a noise floor of around -118 dBm on the 3.1 kHz filter. The accuracy of the receiver looked good when measured against the HP 3335A synthesized level generator, the companion unit for the SLMS in the 1970s. Absolute accuracy was within 0.25db and the attenuator steps within 0.03 dB over an 80 db range.

Although it was then 35 to 40 years since the 3745B was designed, it still impressed with its measurement accuracy and functionality. Its keyboard is a pleasure to use, a tribute to the amount of thought that went into its layout and operation. The instrument’s strength and durability was put to the test on its journey home from Toronto, and it survived admirably! It passed to the National Museum of Scotland along with the other instruments in August 2010 and is now in the permanent collection with accession number NMS T.2010.78.

Comments from Bill Lauchlan in Canada

Bill was a test engineer on the Microwave Link Analyzer production line, transferring to Loveland Division in the mid-1970s as a support engineer, and from there he joined the sales team in Canada

"I was particularly interested in this repatriation project as we did sell one 3745B to Saskatchewan Telecommunications, the telephone company in Saskatchewan. The HP Players in that deal were Hal Dawson - Test and Measurement Field Engineer, Bob Morgan - Technical Computer Field Engineer and myself, Bill Lauchlan - Telecommunications Systems Engineer, while from the customer side we dealt with Don Anderson - Technical Standards Engineer.

We arrived in Regina the night before the Demo and went to a favourite restaurant for supper. Hal Dawson had Baked Alaska for dessert and was "sick as a dog" all night. So much so that he couldn't go see the customer for the Demo. I had to do it all myself. I never did get the commission cheque for that!

Don Anderson was one of these maverick type progressive engineers who was anxious to move ahead without waiting for the big national project which was called TRIMMS[22]. He got his way and later proved to be a great ally of ours. The 3745B was installed at the main toll office (Regina Main) in Regina. At the time this was one of only two Class 1 switching offices in Canada.

This would have been my first desktop-controlled SLMS system installation in Canada and used an HP 9845 controller. When we sent the quotation, we had to avoid the word "computer" because they had to buy "computers" from DEC.

The System had to monitor two sites, Main and a Junction site. 3754A/55A switches and controller were used to access the traffic, however we could not afford to use a second SLMS for the Junction Site, so we came up with a piece of "Scottish innovation" instead!

The 3754A Switches were installed at both locations and the baseband from the remote switch was then "trunked" via an equalized cable link to the 3745B at the main office. We sent the control pulses for remote switch selection by toggling a line on a data modem. So we got two sites with one SLMS. I developed all the software in my "Skunk Works" in Toronto (with some ideas borrowed from the 37050S System).

We later went on to sell similar systems with the same software to Northwest Tel (Whitehorse, Yukon) to monitor Alaska Traffic and also Maritime Tel (Halifax) and Terra Nova Tel (Gander, Newfoundland)."

---

[22] Trans-Canada Remote Interface Measurement and Management System, originally developed by Telecom Canada to monitor circuit quality on the trans-continental links interconnecting international circuits across the Atlantic and Pacific.

6

Chapter Six

Telephone Line Analyzers

Simple measurements on telephone circuits go back to the early days as soon as there was interest in the quality of a connection, particularly on longer distance circuits. The main criteria being tested were the signal level or loss of signal (attenuation) along the circuit and the level of background noise which, in the early days, was usually induced noise from power lines and electric motors etc. as well as crosstalk from other adjacent telephone circuits.

The basic telephone circuit is typically a pair of wires connecting the two telephone handsets through a switch or exchange. The exchange interconnects the different telephones and generates the ringing signal to alert the called subscriber. Although early telephones incorporated their own batteries, it soon became normal for the exchange to supply -48V to the handset through the line so the system was very robust and "always on". The Plain Old Telephone System (POTS) is similar today, although telephones and switches are much more sophisticated 100 years later.

The telephone connection from the handset to the exchange is a pair of copper wires twisted together forming a "balanced" circuit. The twisting and balancing helps to minimise induced common-mode noise and crosstalk mentioned earlier. This is called a 2-wire circuit and carries both directions of transmission. The transmit and receive directions are separated at the handset and exchange by means of a hybrid transformer, providing separate transmit and receive paths creating an equivalent 4-wire circuit. It isn't perfect, so there is always some breakthrough between the two directions, creating a "sidetone" in the telephone ear-piece. This isn't a problem for voice communications and in fact some "sidetone" is desirable acoustically. If higher quality is required, then a separate pair of wires is used in each direction to create a proper 4-wire circuit.

The twisted-pair connection between handset and exchange is referred to as the "local loop" and in some situations can be several miles long. Long circuits cause loss of signal, particularly at higher voice frequencies[1]. This can be mitigated by adding "loading coils" or inductors in series with the line at approximately one mile intervals which create a low-pass filter effect with the inherent line capacitance[2]. The attenuation is then more even across the basic voice bandwidth, although it increases more steeply beyond that.

So, telephone line measurements amounted to some DC voltage and resistance measurements, and the measurement of signal levels and noise over the basic voice frequency band up to around 3.5 kHz. Shorter telephone lines can have a larger potential bandwidth in the local loop, however if the connection is over a multiplex transmission system, often referred to as a carrier section, then the bandwidth has a restriction of 300 to 3400 Hz. The carrier system might be analogue frequency division multiplex (FDM) or more recently digital, using pulse code modulation (PCM). These carrier systems had a particular affect on group delay as discussed later.

HP's first entry into this market was in 1962. The company had a long tradition of making audio signal generators starting with the iconic Model 200A in 1939. By the early 1960s, HP introduced the 204B, a new compact solid-state oscillator and a matching AC voltmeter, the 403B. Both were just over 5" (130 mm) wide and could be operated from internal rechargeable batteries giving 40 hours operation. The new HP 3550A Portable Test Set was a field portable case combining the 204B and 403B with a third module, the 353A, which contained impedance matching transformers and an attenuator. This third module in effect turned some general purpose test equipment into a telephone test set by providing the appropriate balanced line interfaces. It was probably the earliest dedicated telecommunications test set from HP aimed at telephone companies. Developed at the Loveland Division in Colorado, the test set sold for around $1000. It was useful for checking signal levels and line attenuation, but as the voltmeter was wideband, it might have given misleading measurements of noise and crosstalk.

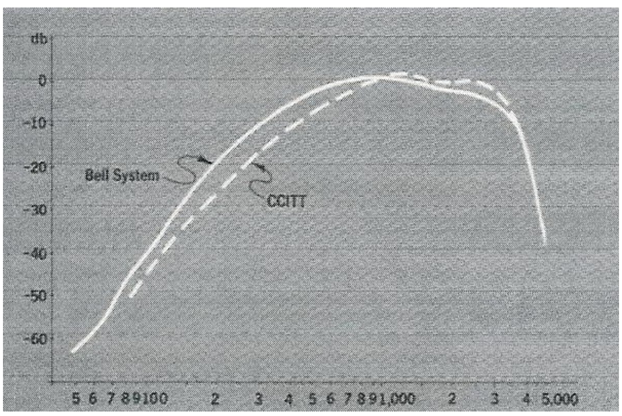

The accurate measurement of noise was addressed by the next product from Loveland Division, the 3556A Psophometer. This curiously named instrument derived its title from the Greek word "*psophos*" meaning noise. It was in effect the 403B AC voltmeter described above with the balanced input of the 353A, but incorporating some bandpass filters specifically for accurate measurement of noise and crosstalk. In particular, it incorporated the "psophometric" weighted filter defined in the CCITT (now ITU-T) recommendation P.53. This filter, and the corresponding C-message filter used in the Bell system (shown here), was intended to give a weighting to the noise spectrum corresponding to the subjective effect in the human ear. The 3556A was introduced in 1970 and sold for around $800.

[1] Resistance in long circuits also reduces the exchange battery voltage reaching the handset so thicker gauge wire has to be used.

[2] The theoretical background for this (the "Telegrapher's Equations") was developed in the 1870s by the British electrical engineer Oliver Heaviside, nephew of Sir Charles Wheatstone one of the inventors of the telegraph. Heaviside was a brilliant mathematician credited with the practical application of Maxwell's Equations and in 1880 with the invention of the coaxial cable. Loading coils were first used on long telephone lines in 1899 by AT&T and were typically 88 mH inductors spaced every 6000 feet.

From early 1972, the Psophometer was manufactured under licence at South Queensferry. On the production line it rapidly acquired the nickname "*Piss-off-ometer*"!

So far as analogue telephone channel measurements were concerned, that would have been the end of the story, except that by the 1960s a dramatic development was taking place that would change everything – the first commercial computer systems were introduced. Soon there was the need to send data between computers and terminals over the communication network. To do this, the computer data was converted to a modulated analogue signal that could be transmitted through the limited bandwidth of the telephone channel. The device that did this was called a modem (modulator and demodulator). Some of the first modems were introduced in the Bell System in the early 1960s and used a simple idea of two audio tones within the telephone bandwidth, one representing "1" and the other, "0". Called Frequency Shift Keying (FSK), these early modems had a throughput of around 300 bits/second. If the telephone channel sounded good for speech, it would probably work with the data, and in fact after setting up the call, one put the handset into an acoustic coupler which allowed transmission of the switched tones. These were the first dial-up modems.

For higher data rates of 1200 or 2400 bits/second, a more sophisticated modem was required using Phase Shift Keying (PSK). This was pretty much "state-of-the-art" until the late-1970s, and these modems required leased or private lines which were 4-wire circuits with carefully specified transmission parameters to ensure the modem signal got through without degradation. Several extra measurements were required in addition to the basic ones discussed earlier. The most important of these was the measurement of phase variation versus frequency across the telephone channel bandwidth, a parameter which went undetected in voice communications but could cause problems with data.

For ideal transmission of the modem signal, the channel had to have a linear phase characteristic or constant (flat) delay for all frequencies in the channel. Delay is the differential coefficient of phase versus frequency and is measured as Group Delay (GD) distortion, or in North America as Envelope Delay Distortion (EDD). When group delay is flat, the modulation on the modem signal (carrying the actual digital information) will be resolved by the receiver without degradation since there is then no time-smearing or interference between the data symbols. Here is the CCITT Group Delay specification.

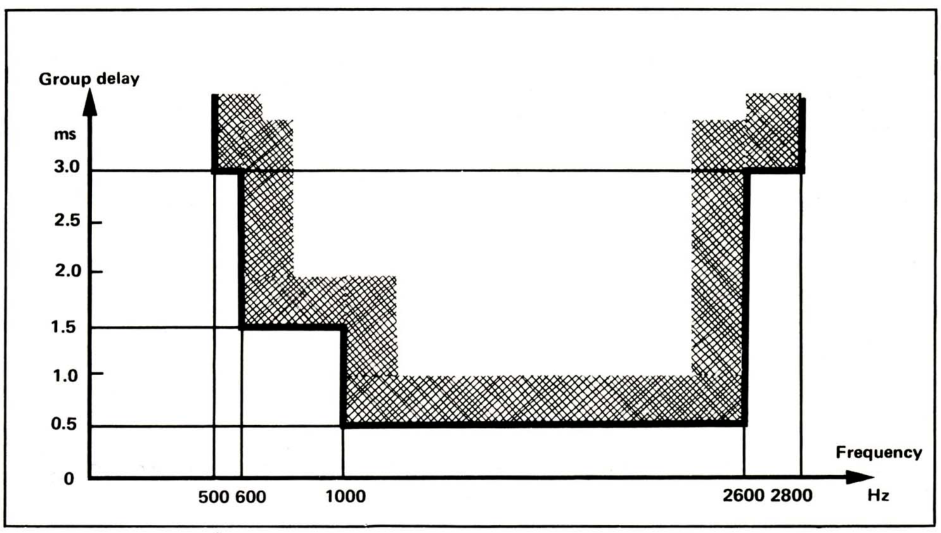

CCITT M.1020 Group Delay Distortion Limits

These complex modem signals were also more vulnerable to another group of impairments called transients. These included impulse noise (random clicks), gain hits, phase hits and dropouts (all sudden changes in the level or phase of the received signal often caused by network switching). Yet another impairment was phase jitter, where the whole modem signal was phase modulated by interference from say power line frequencies. Over the years, modem bit rates increased to 4.8 and 9.6 kb/s by the 1980s. These used even more complex modulation schemes so were more sensitive to impairments and the need for measurement increased[3], although adaptive equalizers for amplitude and group delay non-flatness were introduced to alleviate this problem. The full set of specifications was defined in ITU-T Recommendation M.1020, and the measurements required in M.1060.

Product Numbering

Before embarking on a survey of South Queensferry's telephone line analyzers, a few comments on product numbering are appropriate as there were some anomalies with this product family and some overlap between three HP divisions.

Those familiar with the history of HP instruments will be aware that the first two digits in the product number are associated with a particular HP division. So the early telephone line testers and other signal analyzers from Loveland Division[4] all start with 35. South Queensferry's products start with 37 and this continued until the late 1990s when HP changed to a new numbering convention with a letter prefix.

The third entity involved in this business was the Delcon Division in Mountain View, California. Acquired by HP in 1965, they specialised in ultrasonic testing[5] and in the early 1970s developed various telephone line fault locators. In 1981, this operation moved to Colorado Springs and became the Colorado Telecommunications Division (CTD), specialising in telephone line measurements (mostly for North American markets) and protocol analysis for datacommunications. The prefix for Delcon/CTD products was 49. The significance of all this will be clear a little later.

3770A Amplitude/Delay Distortion Analyzer

At the end of 1971, South Queensferry won a large contract worth $343k to supply 75 Voice Channel Delay Test Sets to British Post Office Telephones[6]. The complete contract was for 150 units, which was split between Queensferry and the German test equipment manufacturer, Wandel und Goltermann (WuG). The equipment had to be delivered by early 1974. Thanks to cooperation between the Division and the sales office in Slough, it was the second largest single order they had ever received. The test set would measure group delay and amplitude flatness across the telephone channel using the method recently defined in the CCITT (now ITU-T) Recommendation O.81. This method had been

---

[3] A useful survey of telephone line impairments is available in Communications Network Test & Measurement Handbook, Chapter 25 "Analog Measurement Instrumentation" by Ronald D. Lowe.
[4] Loveland also used 34XX numbers for its voltmeters and multimeters, and 33XX for its synthesizers and signal generators.
[5] See HP Journal May 1967 http://www.hparchive.com/Journals/HPJ-1967-05.pdf
[6] Like many other post offices round the world at that time, the British Post Office handled both mail and telecommunications services.

adopted from one developed by Wandel und Goltermann[7] and at the time their test set, the rather cumbersome LD-2, was the only game in town. So it was an achievement by the HP team to win this contract considering the Division had yet to design and ship the product! No doubt, in those far-off days the preference was to "Buy British" rather than give the entire order to a German company. Finlay Mackenzie, who was involved in the contract negotiations at the factory, recalled that the deal nearly got them into a lot of trouble. HP didn't usually take orders for equipment that wasn't even designed.

The pressure was on to complete the 3770A project in two years, especially with the alarming discovery that WuG's LD-2 had been superseded by the more competitive LD-3. The design team led by Mario Pazzini had a challenge as how best to implement the measurement method defined in CCITT O.81, the standard used in Europe and internationally. The principle was that the transmitter would send test signals alternately at a reference frequency (usually at 1.8 kHz, roughly the middle of the 3.1 kHz telephone channel) and at the variable frequency under test. The receiver, which was at the far end of the communications link, would lock on to this alternating sequence and measure the relative group delay and amplitude differences between the reference the test components of the composite signal. As there was no other link between the transmitter and receiver, the distant receiver had to work everything out just from the received signal.

Furthermore, in order to make the measurement more automatic, the transmitter would sweep the test frequency up and down the telephone channel bandwidth and the receiver had to follow at the far end and measure the frequency being transmitted. On top of this, the reference and test frequency carrier signals were amplitude modulated at 41.66 Hz, and this modulation signal was used to measure the relative phase delay between the reference and test envelopes at the receiver. To help the receiver lock on to the sequence, a short burst (four cycles) of amplitude modulation at 166.6 Hz preceded the changeover from reference to test carrier frequency. This complex signal needed to be generated with good frequency and amplitude stability and the various frequency limits for the sweep and reference needed to be set from the front panel. An analogue circuit solution looked challenging. The digital approach was even more challenging, given the two-year timescale and the technology available. It was a venture into unknown territory in 1972.

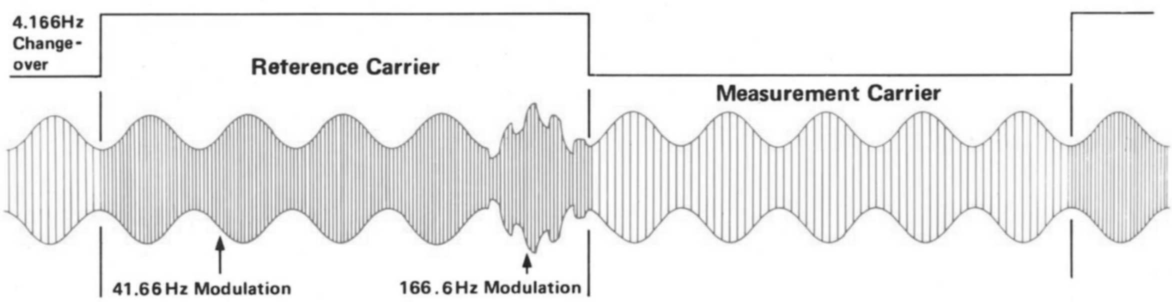

TRANSMITTER O/P SIGNAL

A young design engineer on the project, David Guest, recalls seeing a key paper in the IEEE Transactions on Audio and Electroacoustics[8] in 1971 describing the novel principles of what has now become known as direct digital waveform synthesis. It looked like a

[7] W&G invented the method in the late 1950s and the first product to use the idea was the LD-1.
[8] "A Digital Frequency Synthesizer" by Tierney, Rader and Gold, IEEE TAE March 1971

promising idea which could be implemented with the new logic ICs that came on the market in the late 1960s. In the typical HP spirit of innovation at the time, David was told to get on and design something. It turned out to be the perfect solution for generating the complex test signal shown above.

The 3770A was certainly one of the earliest commercial applications of direct digital frequency synthesis[9]. The technique had several advantages. Firstly, all the frequency settings were tied to a crystal reference clock oscillator of high stability, so drift was eliminated. Secondly, the digital approach made control of the frequency output easy to implement through front panel buttons and thumb-wheel switches which simply had to control logic circuits. Thirdly, because the signal was being generated by Read Only Memory (ROM) look-up, changes in frequency were phase-continuous and so eliminated phase transients. David Guest recalled that the ROM had to be explained to "management", as such an "exotic" device had not been seen in South Queensferry before!

The design used the equivalent of 32,768 (2^{15}) sinewave samples stored in ROM. The ROM was addressed by a digital 15-bit accumulator driven by the clock source. The contents of this accumulator represented the phase of the sinewave, so the faster it was incremented, the higher the output frequency. It was clocked at 327.68 kHz so the smallest frequency increment was 10 Hz and the top carrier frequency set at 20 kHz, giving a minimum of about 16 samples per cycle. The parallel words generated by the ROM fed a digital-to-analogue converter which after filtering produced a remarkably clean sinewave output. The modulating signals were also generated digitally, which allowed inclusion of a very elegant way of introducing a precise one millisecond delay to verify receiver accuracy.

In the 21st century, storing 32,000 sinewave samples would be trivial, but in the early 1970s it required the memory of a medium-sized computer and was quite impractical for a portable measurement instrument. The design team came up with some inventive solutions to reduce the storage requirement. The first, and simple step, was to store only one 90 degree quadrant of the sinewave and generate the other three quadrants by easy binary-arithmetic manipulation of the ROM output word. This gave a four-fold reduction. A substantial further reduction came from an ingenious use of a trigonometrical approximation for the sine function. By making a real time calculation, interpolated values were produced from a much compressed set of stored values. Together, these reduced the storage requirements to what was practical with early logic ICs, and a ROM of only 1024 bits was required.

In the May 1973 issue of the Division newsletter "*Readout*" the 3770A was described as containing 400 logic integrated circuits and representing the most sophisticated logic design yet undertaken, with up to 14 R&D staff working on the project. The 3770A used digital implementations in a number of areas, which avoided the need for calibration adjustments invariably needed in analogue circuits. A bi-directional counter circuit was used to determine the relative group delay between the received reference and measurement cycles, giving a direct readout of the relative time difference. This had the advantage of eliminating any errors in the 41.66 Hz modulating frequency, and also allowed a simple way of averaging out noise by accumulating counter residue over 4 or 16 cycles and dividing the result. A similar up/down counter was used for measuring relative

[9] In contrast to frequency synthesizers that use digital dividers to phase lock a tuneable oscillator to a reference.

amplitude and absolute signal level by means of a dual slope technique frequently used in digital voltmeters. The timing and control logic also determined the start and stop points for measurement to minimise the settling time of filters and where necessary to eliminate the effect of the modulating signal. Yet more circuitry measured the incoming received frequency using a digital frequency counter.

The whole instrument was run by a logic state machine which controlled the measurement hardware, read the push buttons, thumbwheel switches and slide switches on the front panel, and drove the LED[10] displays. The development slightly preceded the first Intel microprocessors, so the whole instrument control was implemented in logic hardware rather than firmware. Interestingly, Ralph Hodgson designed the control logic using principles taught by Bernard Howard on the Digital Techniques M.Sc. course at the local Heriot-Watt University in Edinburgh, which Ralph and a number of other South Queensferry engineers took on a part-time basis. Howard enthusiastically promoted a more organised and rigorous approach to the design of logic state machines, which avoided many of the pitfalls of early digital design.

Although housed in a single box, the transmit and receive sides of the instrument were completely independent. The transmitter could be sending a test signal in one direction while the receiver simultaneously performed measurements on a quite different signal from a distant test set, possibly one made by another manufacturer!

The 3770A could drive an X-Y plotter to give a graphical presentation of the channel amplitude and group delay characteristic, and special graph paper was supplied pre-printed with the M.1020 masks, as shown above.

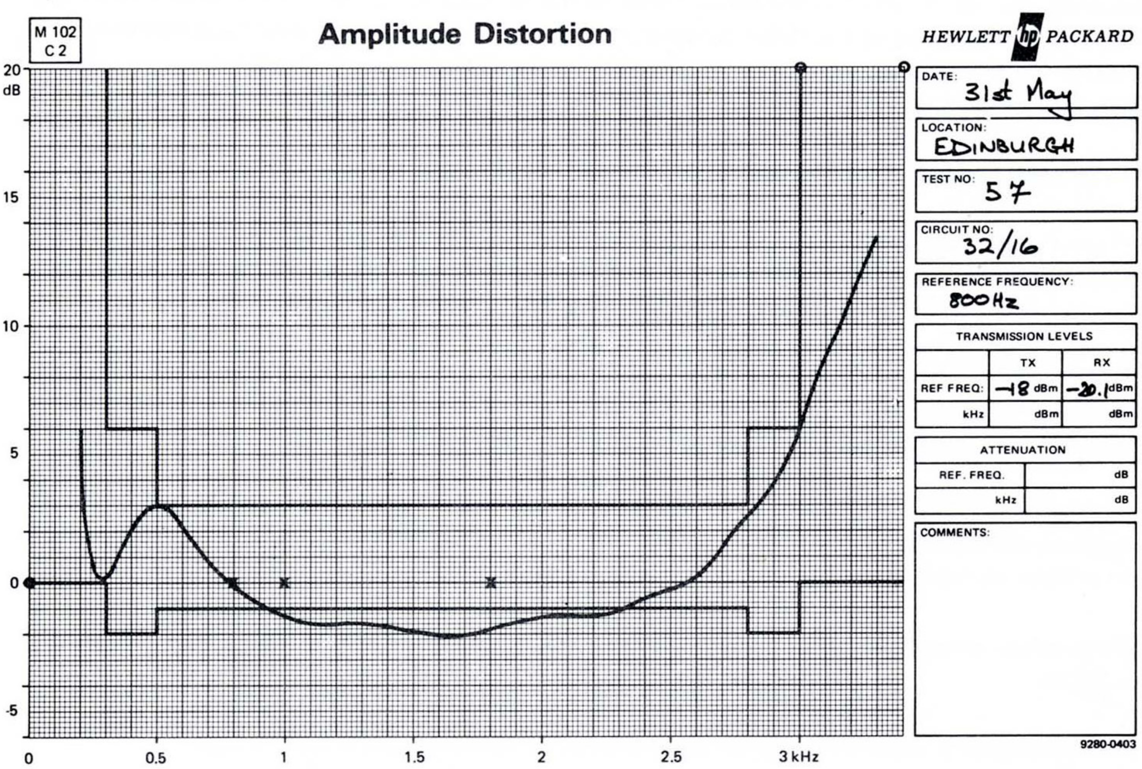

[10] Light Emitting Diodes. HP had just introduced the first 7-segment displays which were used in the 3770A

A couple of further innovations rounded off this remarkable instrument. It was the first South Queensferry product to use a fully portable package. This was borrowed from the new HP 1700 series oscilloscopes. The extensive use of digital hardware did have a drawback: those early digital logic ICs were power-hungry.

Using a conventional analogue power supply would have required a heavy mains transformer, compromising portability. The problem was solved by the design of a switching power supply which was 75% efficient and used much lighter components. Apparently, the lab benches were littered with burnt-out devices for quite some time while this difficult circuit was being perfected. Designed by Rajni Patel, it was one of the first switching power supplies in any HP instrument. David Guest recalls an informal lab visit by Bill Hewlett (HP co-founder), whose interest in the project tempted Rajni to exaggerate the supply's efficiency:

"*From Hewlett's momentary change of expression it was clear he could detect that something was amiss. With a couple of well directed questions, put with the utmost tact and a characteristically modest air, he elicited the response that perhaps the efficiency was rather less than first claimed!*"

A more detailed description of the instrument and application (written by David Guest who took over as Project Leader when Mario Pazzini returned to Italy in April 1973) is published in the *HP Journal* of November 1974[11].

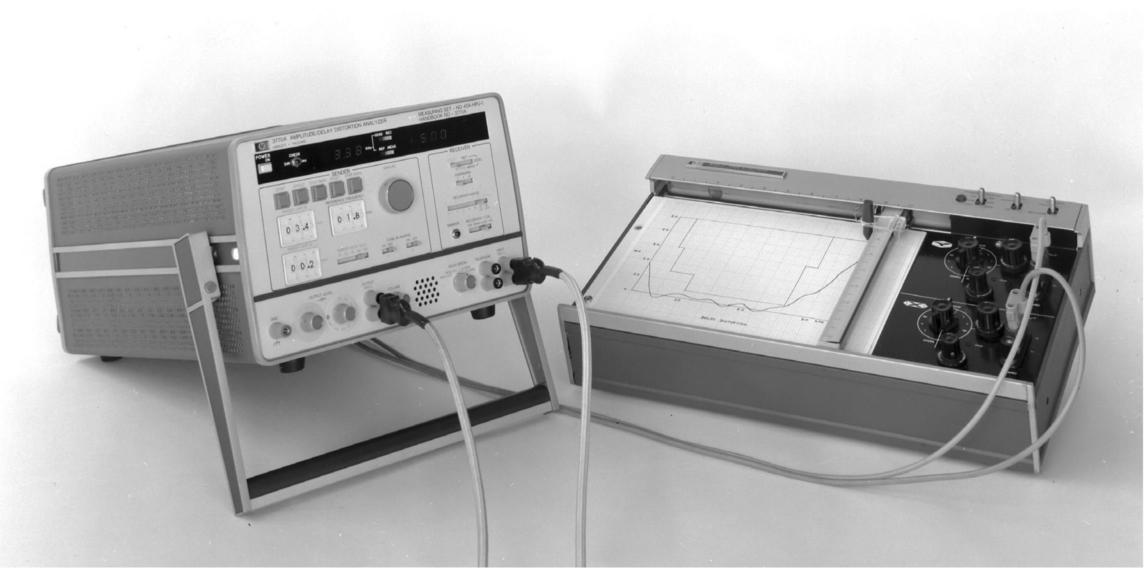

Following some prototype visits to key customers, the 3770A was introduced in the summer of 1974, close to target, although early production runs were committed to the British Post Office order. The unit cost $6k to $7k and the average sales volume was 80 units per year until 1977 when the 3770B was introduced. A total of 430 units were shipped and total revenue was approximately $3.5M. The Post Office order set a trend, as more than half the total units were sold in the UK, while only one was sold in Germany – possibly Wandel und Goltermann checking out the competition.

[11] "Simplified Data-Transmission Channel Measurements" by David H. Guest, HP Journal November 1974 pp 15-24. http://www.hparchive.com/Journals/HPJ-1974-11.pdf

David Guest remembers an amusing incident when, after showing the new instrument to the British Post Office team, they asked them what they thought, no doubt hoping for compliments about the novelty and precision of South Queensferry's digital implementation. However, the bland response was simply, "*To the Post Office, something is either 'satisfactory' or 'unsatisfactory'*".

The digital implementation did have an unexpected benefit however. South Queensferry anticipated that a licence would be required from Wandel und Goltermann to use the two patents they had on the measurement method, and negotiations were started by HP's patent agent, Knud Schulte. However, it turned out the 3770's innovative digital implementation had avoided the patent infringement, so WuG were informed, to their surprise, that no royalties were forthcoming! Instead the South Queensferry team were able to file for patents based on their own designs (US Patent 3970926 and UK Patent 1429617, authored by Peter Rigby). Here is the 3770 team in 1974, David Guest front left.

3770B Telephone Line Analyzer (Unit Preserved)

Soon after the 3770A was launched, work started on the B model, led by Rajni Patel (later Kansagra). The A model was designed principally for measuring group delay and amplitude flatness to ensure the telephone channel met the mask limits defined in M.1020. If the operator moved a rather inconspicuous slide switch on the back panel, the modulation could be turned off and the 3770A could then be used as a signal generator and level detector up to 20 kHz. For the other measurements specified in M.1060, the user needed to have at least one additional test set. The objective of the B model was to put all the necessary measurements under one handle. Reflecting this, the 3770B was renamed as the Telephone Line Analyzer.

The new measurements added were psophometrically weighted noise and impulse noise. The customer could also select optionally either noise with tone or crosstalk measurement. The "noise with tone" measurement was the psophometric filter with a narrow notch at either 820 Hz or 1020 Hz. This holding tone was sent from the transmit end and then removed by the notch filter in the receiver. By measuring the level of the tone, a direct measurement of signal to noise ratio was calculated. The holding tone was necessary in some systems for proper operation, for example in a PCM system the tone is needed to activate the digitization process. The new instrument also had a slave facility, which used the return channel to transmit the measurement results from the far-end back for display on the local 3770B. To make more accurate noise measurements, a true RMS detector was incorporated.

The 3770B Telephone Line Analyzer was launched in 1976 and sold for roughly $2.5k more than the A model. Overall, although more expensive, it was a better product as in many cases it was the only instrument the engineer needed to qualify a line for data transmission. It largely replaced the 3770A which was discontinued in the late 1970s. Production volume was around 10 units/month, although it fluctuated significantly due to some further large contracts in the UK.

The British Post Office was clearly enthusiastic about the 3770A and B, and the B model was designated P.O. Tester No.56A. Other UK companies such as utilities and the railways followed suit and a total of 439 B models were sold in the UK. Germany bought four units. As with the earlier model, UK sales dominated the overall total of 1200 units. Production ceased in 1986, with total revenue of over $11M. If the UK market penetration had been replicated elsewhere in Europe and internationally, the sales could have run into several thousand units.

The 3770A/B was one of several highly innovative products developed at South Queensferry in the 1970s, which broke new ground and pushed the limits of technology at the time. Its legacy was that some of the design concepts were used in the PCM testers developed a few years later (see Chapter 8). The preserved 3770B is a P.O. Tester 56A version, manufactured in early 1981.

3771A/B Data Line Analyzer

Although the 3770B satisfied the measurement needs on data circuits for 90% of applications, some customers wanted to check the full range of measurements defined in M.1060. The additional measurements included further transient impairments such gain and phase hits, dropouts and three-level impulse noise (three different thresholds were used for counting impulses rather than just one in the 3770B). Another measurement not covered in the 3770B was phase jitter.

A new instrument, introduced in 1978, filled in these gaps and was a companion instrument to the 3770A/B. The 3771 used the same portable package and had a similar user interface. It was mainly a measurement instrument, though it could send a couple of fixed tones required for some measurements. Its main claim to fame was that it made all the transient measurements in parallel. Since these transient events were random, a measurement period of at least 15 minutes was required to accumulate results. Doing all the measurements simultaneously saved time. The new product also measured phase jitter, frequency shift as well as weighted noise (psophometric or C-message).

Two versions were produced. The A model followed the CCITT specification M.1060, while the B version worked to the North American Bell specification defined in Bell Publication 41009.

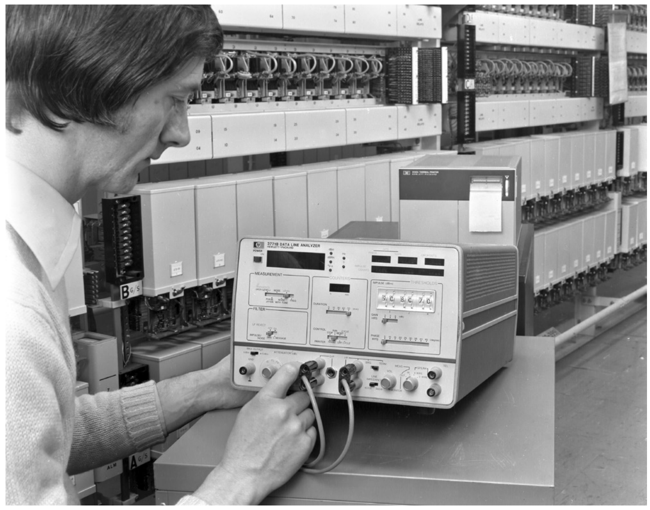

Neither model was a big seller. Perhaps relatively few customers needed to do a comprehensive test of all the specifications which was a time-consuming business. The price was around $8.5k. The 3771A sold 435 units, fairly well spread across the global market, but again with a disproportionate fraction (25%) in the UK. Sales of the 3771B to the North American standard were very disappointing at 87 units. Perhaps this wasn't surprising given the high level of internal competition from the US divisions, discussed in the next section. Both units ceased production in the mid-1980s along with the 3770B.

A Crowded Market

While South Queensferry was developing these products, the other two HP divisions in this market were busy with their own development programmes. Loveland Division in Colorado had been the first to enter in 1962 with the 3550A combo-instrument described earlier. This was followed by the 3556A Psophometer in 1970 and a couple of years later a North American version of this, by the 3555B, which made noise measurements using the C-message filter.

In the 1975 HP Catalogue, Loveland launched two new products, the 3551A and 3552A Transmission Test Sets, the first for North American standards and the second for CCITT. These well thought-out designs combined the functions of the 3550A (signal generation and level measurement) with the Psophometer, in a neat portable instrument that could run off mains power or internal batteries. They were straightforward instruments, rugged and simple to use and made the basic telephone line measurements the telephone engineer (or in US parlance the craftsperson or "craft") needed. Ideal to have in the back of a truck as the units came with a robust front-panel cover. These simple products were enormously successful and were in production until the late 1980s (some built at South Queensferry), although these were the last telephone line testers from Loveland.

Over in Mountain View, California, the Delcon Division was on a different track. Their new product was the 4940A Transmission Impairments Measuring Set or TIMS for short. It was the first time HP used this title, and was no doubt considered preferable to Transmission Impairments Test Set as the acronym might have caused embarrassment. The 4940A was also new in the 1975 Catalogue. It was the opposite of the 3551/2, since it pretty well included every measurement that could possibly be made on a telephone circuit in North America[12]. It was the "Big Daddy" of telephone line analyzers, being the size of a small suitcase with a front panel 18.5" (470 mm) square and 13" (330 mm) deep. The

[12] See HP Journal August 1974 gives useful details of TIMS measurements and 4940A, and has a good description of the P/AR test. http://www.hparchive.com/Journals/HPJ-1974-08.pdf

catalogue entry portrays it as a friendly beast with a pet name of TIMS – "*with TIMS you can do this, TIMS can measure that*". While the 3551 was ideal for the travelling telephone engineer, you might think twice about dragging TIMS out of the truck. There was quite a contrast in the price too, as TIMS cost $8.5k to $10k in 1975 whereas the 3551 was just $1.75k.

TIMS was the first HP test set to offer the North American Envelope Delay Distortion (the equivalent of the 3770 group delay) and also the P/AR measurement. This stands for Peak to Average Ratio, and was a composite benchmark measurement devised by Bell Labs. It involved sending a comb of 16 fixed tones (generated by shaped pulses) which attempted to simulate the spread spectrum of a live data modem. The receiver simply measured the peak and average values to compute P/AR. A value close to 100 was good. It was found that P/AR was correlated to parameters such as group delay, amplitude flatness and non-linear distortion, but was not an analytical measurement. It was useful as a "go/nogo" test or for benchmarking circuit performance over time and was quick and simple to carry out.

Maybe the "Big Daddy" TIMS was a bit too much of a good thing, as a year or two later Delcon introduced the 4942A which was subsequently replaced by the 4943A and 4944A. These units were all built into a portable HP 1700 case similar to the 3770, and had various subsets of the full TIMS menu and at a more competitive price.

In 1981, by which time the division had moved to Colorado Springs and renamed itself the Colorado Telecommunications Division (CTD), two new portable TIMS products were introduced. These were the 4935A for North America, and the 4936A for CCITT standards. Built into rugged polycarbonate enclosures and weighing just 14 lbs with internal batteries, they clearly targeted the "craft" segment of the market as did the earlier 3551/2. However, the new instruments were even more compact and included the additional measurements of 3-level impulse noise and P/AR in the 4935A. They sold for $3.5k to $4k.

This family of products were now referred to generically as "TIMS boxes". In the 1982 HP catalogue, there were 14 different products addressing this market, including three from South Queensferry. Amazingly, the original 3550 combo-instrument from 1962 and the old psophometers were still listed. Maybe their longevity was down to US Army procedures specifying them for eternity ("Tester, Communications System for the use of"). Overall, it was a somewhat over-populated portfolio.

There was the natural delineation between Bell and CCITT standards, and there was a high-end box (4940A) and low-end craft boxes providing the basic measurements of level and noise (3551A/52A). In between there was a whole variety of products with different mixtures of the telephone line measurements aimed at different users, grades of telephone line, or whatever.

I remember one of the South Queensferry product marketing managers, Geoff Thow, talking about "High-end TIMS, Low-end TIMS and Middle-end TIMS". MIDDLE-END TIMS? I was never sure if this apparent oxymoron was a joke or whether it was serious. There was also Upper-middle-end TIMS and Lower-middle-end TIMS. It was a confusing landscape, but obviously lucrative for HP. Some of the lower-end craft boxes probably had volumes of 200 to 300 a month or more in the early 1980s, and the sector could well have produced $30M to $50M per year for the Company.

4948A In-Service TIMS

Confronted with such a congested portfolio, the Scottish Division went off at a tangent and came up with something completely different.

In the late 1970s, South Queensferry got interested in high-speed data modems at 4.8 kb/s and 9.6 kb/s following the success of the HP-IB Extender products. For various reasons (see Chapter 12) these were not a commercial success, but several design engineers became very conversant with the design of adaptive equalizers and digital signal processing. The idea of the adaptive equalizers was that they automatically compensated for line impairments (for example amplitude and group delay non-flatness) and thereby provided better modem performance. Conversely the optimised settings arrived at by the adaptive equalizers gave a measure of the degradations on the telephone circuit, assuming the originating modem signal was of good quality.

This was the basis of the In-Service TIMS or ITIMS, which would make its measurements on live modem traffic rather than test signals. It was a clever idea that surely could only have been invented at the "University of South Queensferry"! It wasn't just a clever idea though, as there was always strong market interest in assessing network performance without taking systems out of traffic. The unique ITIMS proposal seemed to have a lot of potential. Apart from the advantages of making measurements without interrupting service, the ITIMS was ideal for long term monitoring of circuits, collecting transient impairments or identifying periods of degraded performance. Making either of these measurements conventionally was unattractive because of traffic down-time. The product marketing staff envisaged customers with large data centres using an access switch to monitor many different modem links non-intrusively. It seemed like a good proposition.

Around 1982, a project team was set up and led by Gordon Rhind who had been a senior design engineer on the high-speed modem development. Over time, the project had up to 10 design engineers working on the new product. This was not surprising as ITIMS was one of the most theoretically complex products ever designed at South Queensferry. There was no blueprint for such a product in the market, so research was required in product marketing and R&D to define the capabilities and develop the signal-processing algorithms largely from scratch.

As mentioned earlier, higher-speed modems used combinations of phase and amplitude modulation on a fixed carrier frequency (usually 1700 Hz or 1800 Hz) to transmit digital data through the telephone channel. The idea was that the modem used particular combinations of phase and amplitude, referred to as phase states, to transmit information. Provided the signal arrived intact, the modem receiver could decide which of the phase states had been transmitted from the far end and could recover the encoded data. The phase states for a CCITT V.29 modem operating at 9.6 kb/s is shown here as an example.

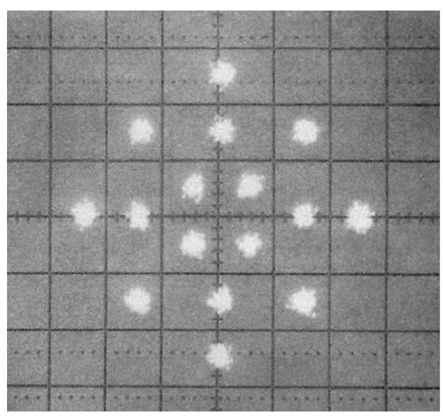

Each "star" in this "constellation" is a "symbol" and represents a unique 4-bit data pattern. So at the sending end, data is received serially in 4-bit segments and, depending on the pattern, is encoded as one of these 16 possible phase states. At the receiving modem the phase state is identified and the 4-bit pattern recovered. This serial to parallel conversion results in the 9.6 kb/s data stream being reduced to a rate of 2400 symbols per second This can be transmitted through the 3.1 kHz telephone channel bandwidth.

Ideally, the phase states in the constellation should be single points but due to noise and transmission impairments, they appear as clusters round each ideal point. The idea behind the ITIMS, was to quantify these variations from the ideal position and interpret the results in terms of the standard telephone line transmission impairments defined in M1020 and M.1060.

Some impairments, such as transients (impulse noise, gain and phase hits and dropouts), show themselves as sudden changes in the mean position of the phase states, while others such as phase jitter cause continuous cyclical changes in position. More complex are the effects of non-flat group delay and amplitude response. These transmission impairments cause "intersymbol interference" which means the residue of previous symbols (or phase states) are superimposed on the current received phase state. On the constellation shown above, this looks like noise and is visually indistinguishable from random noise. However, unlike random noise, this spreading of the constellation points is pattern-dependant (pseudo random) so can be quantified and separated out.

This was a key capability of the ITIMS. Like a high-speed modem, it had an adaptive equalizer to correct the effects of group delay and amplitude flatness in the telephone channel so these effects could be eliminated at the point when the ITIMS decoded the transmitted phase state. It then used this information to drive a second adaptive filter which was in effect modelling the line. This was a transversal digital filter with 128 coefficients. These adaptive equalizers were operating in the time domain and were a series of "registers" each delaying the signal by one symbol time, and then adding a fraction of this delayed signal to the current symbol. This represented the impulse response of the telephone line in the time domain. Being adaptive meant the signal processor went through an optimisation routine, altering the filter coefficients to minimise the error at the output. Since the group delay and amplitude response were static parameters, the adaptive filter would converge on a fixed set of optimised coefficients, whereas the random noise component would be gradually averaged out.

The time domain impulse response defined by the coefficients then had to be converted to the frequency domain to create the familiar group delay and amplitude response plots versus frequency, defined in M.1020 and measured by the 3770. This required a Fourier Transform of the time domain measurements.

Further signal processing and adaptive filtering was used to isolate parameters such as phase jitter, while a digital phase-lock loop quantified any frequency offset from the nominal carrier frequency (e.g. 1700 or 1800 Hz). This was very important as any frequency offset had to be eliminated to stabilise the constellation pattern, which otherwise would rotate in one direction or the other due to accumulating or decrementing phase.

In some ways, making measurements on the actual live modem signal should be more representative of how the modem will be affected by line impairments as the "test" signal occupies the whole channel unlike a single tone. "*The ITIMS sees the channel like a modem sees it*" was the sales message. If the impairments could be measured by the

ITIMS, then they would probably affect modem performance. One area measurement where the ITIMS was potentially less accurate than the traditional tone-based measurements was transients. As mentioned earlier, these produce sudden changes in the mean position of the phase states, so at first sight the ITIMS method would seem to be just what was required.

Take another look at the constellation for the V.29 modem shown above. This picture is a bit misleading as it doesn't really exist in real time. Firstly, it only shows the demodulated modem signal at the instant in time when the receiver samples the signal for phase and amplitude. The modem needs to recover a synchronizing clock in order to sample or "slice" the signal at the right moment. Secondly, at any particular sampling instant, the signal can only take one of the 16 phase-states shown in the picture, so the photograph represents several hundred sampling instants laid on top of each other. Between sampling instants, the signal needed to make a transition from one phase state to one of the 15 other phase states. That meant that in terms of making measurements in the ITIMS, there was a dead zone during this transition period which for 2400 symbols per second is about 400 microseconds. This was a fundamental drawback of the ITIMS compared to the conventional measurement, and it could potentially miss short transient events.

In terms of hardware, the ITIMS did virtually all its processing in digital form using various algorithms running on a specifically designed signal processor capable of executing the calculations in real-time. The input signal was sampled and digitized using a 12-bit analogue-to-digital converter. A high proportion of the routines required multiplication, so the execution unit was designed around a high-speed 16-bit by 16-bit multiplier chip. The design team also looked for ways to minimise execution time through direct memory access and pipe-lining. David Grieve, who worked on the software development, recalled that there was also a secondary bit-slice processor to recover the actual transmitted data.

"*We never introduced this as it was regarded as 'too hot to handle' from a security viewpoint. Needless to say, after the ITIMS was released we got an enquiry from "Maryland Procurement" (US Department of Defense) as to whether a data option was possible!*"

The ITIMS could analyse all the regular modem standards in North America and internationally at 2.4, 4.8 and 9.6 kb/s (about 12 in all) and a year or two after launch, the CCITT V.33 standard at 14.4 kb/s was added as an option costing about $1k. The user could specify the modem standard being used at setup, or the ITIMS would search for it automatically.

A detailed description of the signal processing in the ITIMS is well beyond the scope of this chapter, however much more information is available in the *HP Journal* article about the ITIMS in the October 1987[13] issue.

The ITIMS was housed in a regular HP instrument case (rather than a field-portable package) and was the first of a new range of instruments with a menu-driven interface using a four-line LCD alpha-numeric display. It required relatively few buttons to navigate

[13] "In-Service Transmission Impairment Testing of Voice-Frequency Data Circuits" by Carder, Dunn, Elliott, Grieve and Rhind, HP Journal October 1987. pp. 4 - 16 http://www.hparchive.com/Journals/HPJ-1987-10.pdf

through the menu and was markedly different to many of the other telephone line analyzers with one-button-per-function. It could tabulate the in-service measurements on the display and could send them to an attached printer to provide data logging of results over time. The printer also allowed a plot of group delay and amplitude response against the M.1020 specifications. The ITIMS could also operate out-of-service by transmitting a standard modem signal or tones. There were X-Y outputs on the back for connection to an oscilloscope to display the constellation.

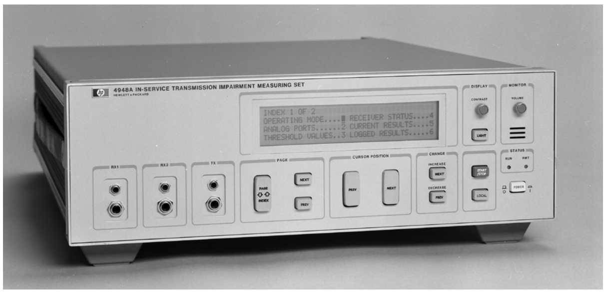

The 4948A In-Service TIMS was launched at the end of 1986. Readers may well wonder why it has a number using the 49 prefix reserved for Colorado Telecom Division (CTD), discussed earlier. The ITIMS did originally have the product number 3774A[14], however before launch there was pressure from the US field to give it a 49XX number so that it appeared to be part of the same family as the rest of the more recent TIMS boxes and protocol analysers. What's in a number? The Division went along with this, no doubt aware that some work would be required to help the sales force sell this unusual product. The unconventional measurement method would inevitably raise questions from customers which the sales engineer would have to address. It would be so much easier to sell another truck load of low-end TIMS, rather than discuss transversal digital filters and Fourier transforms!

At launch, the 4948A sold for around $12.7k, at the top end of the TIMS price range. South Queensferry obviously wanted to recover the large amount of design cost they had invested and felt this unique product could command a premium price. Anyway, the old method of pricing a product based on manufacturing cost didn't really apply anymore to instruments with large amounts of firmware. Sales went reasonably well, with 190 units in the first year and 274 in 1988.

The main selling point of the ITIMS was that it could monitor data circuits non-intrusively allowing customers to adopt preventative maintenance without taking data networks out of service. The Division developed some computer software running under HP's OpenView Network Management System. It controlled the ITIMS and the 3777A channel access switch, so the system could automatically monitor multiple lines. The 37480A OpenView Data Line Monitor (yes, this one had a proper Scottish product number) came out in 1989. Surprisingly, it was not a success and only sold 4 systems, so it was withdrawn a year later. Maybe this was a bad omen, or did it just have the "wrong" product number ….

By the early 1990s, sales of ITIMS were already tailing off, with 88 units in 1992 and only 31 in 1993. By then, the price was around $13.7k, plus $1k if you wanted the V.33 14.4 kb/s modem standard. By 1994, sales had dropped to less than a tenth of what they had been at launch and the product was withdrawn. A total of 1120 units were sold, bringing in about $15M. It sold worldwide with the largest share in North America.

Why did it have such a short life-cycle?

[14] It was also initially called the Non-Intrusive Communications Analyzer (NICA) but this was changed along with the product number.

There was always a question whether the sales force would or could sell the instrument with its unconventional measurement method, however sales did go well early on. There was always plenty of support and literature from the Division to assist the sales engineer, and the basic selling points were good. Perhaps of more concern was the measurement method, which didn't conform to the standard laid down in CCITT M.1060. There was always an issue for customers about the validity of measurements. The whole intention of the ITIMS was that it would be used to monitor critical data links. These modem connections were over analogue leased lines, which were permanent non-switched connections from one site to another. The telephone company guaranteed the line would meet the CCITT M.1020 specification (or the equivalent North American standard). These were also referred to as the "tariffed parameters". The customer paid a high price for these circuits, particularly on international connections. If there was a dispute about performance, the measurements had to be made using a conventional test set.

Obviously these issues were there from day one, but business had been reasonable in the early years. The main reason for the sharp decline in sales in the early 1990s was probably the rapid move to digital services. All the premium data traffic that the ITIMS was designed to service, transferred to digital leased lines either primary rate (1.5/2 Mb/s) or digital services such as DDS[15]. The fate of the ITIMS by its very nature was totally tied into the analogue leased-line high-speed modem business. There was enormous growth and development of more sophisticated and faster modems in the 1990s, but these were largely aimed at the dial-up market for booming internet access. The ITIMS had very little application in switched circuits. Ironically, the basic TIMS was still used in the digital era simply to check the quality of the local loop to support digital services. It needed to make measurements up to 110 kHz, far higher than the 4 kHz top frequency of ITIMS!

However, the short lifecycle and moderate sales don't in any way detract from the remarkable technical achievement of the ITIMS design. Before it entered the market, many people would have thought such an instrument was impossible. Recalling times nearly 30 years earlier, Gordon Rhind commented,

"*Did we really do that? I cannot now imagine how we did those things. It was great fun at the time – textbook stuff coming to life in an absolutely vivid way.*"

The ITIMS had a short window of opportunity: it couldn't have been designed before the 1980s because high-speed digital signal processors weren't available, but the same digital revolution curtailed its potential market in the early 1990s.

4947A Transmission Impairment Measuring Set (MacTIMS)

After the complexities of the ITIMS, the reader will be relieved that the remainder of this chapter returns to more familiar territory. In the early 1980s, South Queensferry developed a new instrument for testing Pulse Code Modulation (PCM) terminals (see Chapter 8 on PCM testers). This new instrument used an in-house digital filter chip to replace the many analogue filters required in PCM test. It also had the option of added data line (TIMS) measurements. In the 1985 HP Catalogue, the 3776A/B PCM Terminal Test Set also put

[15] Digital Data Service in North America provided rates up to 56 kb/s over the local loop. In the UK, the equivalent was Kilostream. High-speed data modems also moved up to 19.2 kb/s which ITIMS didn't test.

in an appearance in the Datacommunications section offering a wide range of TIMS measurements with added PCM.

The Division then had the idea of producing a TIMS-only version of this test set by eliminating the PCM part. This became the 4947A Transmission Impairment Measuring Set. Like the ITIMS, it started life as the 3773A and became the 4947A at launch to fit in with the portfolio.

Based on the 3776, the front panel was redesigned for the TIMS measurements and several new features and interfaces added for the telephone line market. These included DTMF dialling and the ability to run a sequence of measurements automatically and store results end-to-end. It had good ergonomics and referenced the popular 4935A keyboard arrangement. It featured a fairly complete set of TIMS measurements for the North American market.

When it was introduced in mid-1986, there was already a new high-end TIMS product, the 4945A, from Colorado with a complete set of TIMS measurements. It had replaced the 4940A "Big Daddy" TIMS from ten years earlier. The new product used a CRT menu-driven display for control and measurement results, a first in HP's family of TIMS boxes. Side by side, the 4947A and 4945A had very similar feature sets. The main difference was that the 4945A went to 110 kHz whereas the 4947A was limited to 5 kHz by the digital filter technology. The 4945A cost nearly twice as much as the 4947A which had a list price of $8k. The 4947A looked like a high-end TIMS but with a "middle-end TIMS" price. Both instruments were designed for the North American Bell standards.

To promote this new entry into HP's well-populated TIMS portfolio, the Division launched it to the field sales force as the "MacTIMS", its Scottish personality being promoted with photos of the instrument draped in tartan and bagpipes.

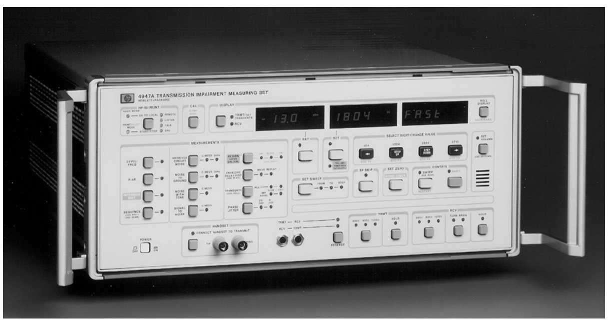

Despite its comprehensive feature set and competitive price, it didn't sell very well. The most significant difference with the 4945A was the top measurement frequency. Like the 4945A, it looked more like a bench-type instrument rather than field portable, and perhaps for that application the top frequency limit was a problem.

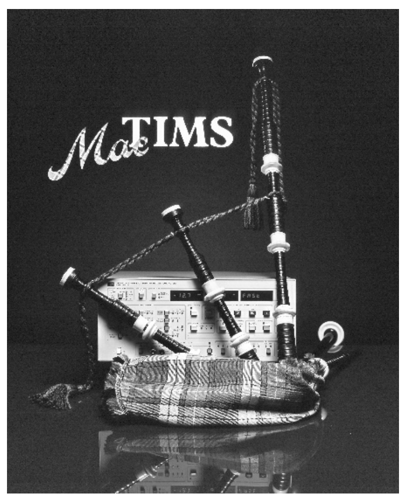

Unit sales peaked at around 13/month in the late 1980s, a miniscule volume in a total market that probably reached around 1000 units a month including low-end TIMS. By the early 1990s, as the market went digital, the volume dropped to 3/month, presumably the 5 kHz top measurement frequency then becoming more of an issue.

The "MacTIMS" met its end in 1994. Total sales of 770 units (almost entirely in North America) amounted to $6M in revenue.

4934A Transmission Impairments Measurement Set (Unit Preserved NMS T.2010.71)

In the late 1980s, South Queensferry took over product line responsibility for all the TIMS instruments although some were still manufactured at Colorado Telecom Division. Thereafter, CTD specialized exclusively in protocol analysis for LAN and WAN systems. For some time the 3551/2 and 4935A TIMS boxes had been manufactured in Scotland, and some of the earliest transferred products manufactured at the new Queensferry Microwave Operation in 1984 were these instruments.

Product line responsibility meant that the Division records contain a fascinating insight into the business being done by HP in this market in the late 1980s:

HP Product	1987		1988	
	Units	Revenue ($)	Units	Revenue ($)
3551A ($3.6k)	475	1.7M	341	1.2M
3552A ($4.2k)	553	2.3M	493	2M
4935A/S ($4k)	2870	11.5M	2040	8.1M
4936A ($4.6k)	466	2.1M	490	2.2M
4937A/S ($5k)	260	1.3M	220	1.1M
4945A ($15.3k)	231	3.5M	172	2.6M
4947A ($8k)	159	1.3M	152	1.2M
4948A ($12.7k)	186	2.4M	274	3.5M
Totals	**5200**	**$26M**	**4200**	**$22M**

From the above, it is clear that the 4935 TIMS, introduced by CTD in 1981, was a high volume product and very popular with the North American customers. With such a large customer base, HP was aware that the 4935 had some weaknesses. The polycarbonate enclosure and particularly the plastic handle were not well liked. Battery life was poor and the batteries were unreliable. In terms of measurement, the customer had to choose between P/AR and noise-to-ground, whereas they would have liked both. Alignment and calibration in production were largely manual as there was no remote control on the product, and this pushed up production cost.

In what seemed a rather unusual move, corporate management in Palo Alto asked South Queensferry to develop a replacement for the 4935. Their interest in this, apart from the obvious sales volume, was possibly that the corporate R&D manager was Bob Allen who had earlier been R&D and general manager at Delcon at the time of the 4940A “Big Daddy” TIMS in 1975. Perhaps Bob had convinced his corporate colleagues that this old analogue communications product was worth one more iteration.

This didn’t seem to go down too well at the Division. Robin Myles, the R&D Manager, recalls, “*Having just about moved the entire lab over to digital communications products, we were ‘told’ by Group to take on development of the next generation TIMS for the phone companies in North America.*” Since the 4935 was so well established and written into procedures, the new product needed to be as close to the old instrument as possible in terms of specification and user interface.

Following instructions, a project team was organised and led by Andy Batham who reported to section manager Gordon Rhind, previously the project leader on the famous ITIMS product. South Queensferry had a penchant and reputation for technically advanced solutions (the R&D approach) and was sometimes accused of "over-egging the pudding". The new development would reinvent the old 4935 using digital filters and digital signal processing (DSP) along the lines of the 4947A and 4948A. Mike Kerr, from product marketing, recalled there was an early "concept" model with a hinge-up keyboard and LCD display. What Corporate and the US sales force probably had in mind was a revamp of the old analogue instrument.

Inevitably politics ensued. At the time, the Division had a North American Sales Manager, Tom Smith, based in Colorado Springs. From previous experience, he was wary of South Queensferry's "high-tech" solutions and extended development times. He and the US field fed back their concerns and some sales managers probably lobbied Corporate, who in turn pestered the Division Manager, Finlay Mackenzie, telling him the job needed to be done quickly. Meanwhile, the Division's marketing team, reflecting the US field position, advised cancelling the project as the return on investment would be poor.

All this negativity had a devastating effect on the morale of the project team. After nearly a year of uncertainty, it all came to a head one Friday afternoon when Gordon Rhind went to Robin Myles and said the team had no desire to carry on. Finlay called an impromptu meeting with all the key players in his office. After several hours of deliberation, it was decided to cancel the project due to lack of support.

That weekend in September 1987, Finlay and Robin Myles were due to fly to California for a Division Review meeting. On the Monday evening they had dinner (Italian apparently in Rohnert Park, Santa Rosa) with the Group Marketing Manager, Gil Reeser, who left them in no doubt they had made the "wrong" decision. Next, it was a phone call back to the factory informing the team that the project was reinstated – and make it snappy. It was a classic case of a popular acronym at South Queensferry – JFDI – (Just F*****g Do It!).

Around this time, or shortly afterwards, the DSP approach was abandoned and the team focused on simply redesigning and improving the existing 4935 circuitry to create the new 4934A TIMS. To use another familiar Queensferry acronym, it became a Q-TAP (Quick Turnaround Project). It did mean that nearly a year's work for some of the team was shelved, including work on a new processor and firmware.

Alex Ballantyne, a senior design engineer on the project, recalled some of the work that was done:

"*We implemented the direct digital synthesizer on an ASIC (Application Specific Integrated Circuit), replacing the discrete logic in the 4935 and saving space. The receiver filters were done with a switched capacitor active filter removing the need for adjustment, while the notch filter used for noise-with-tone measurements was redesigned as a thick film hybrid, laser-trimmed by the manufacturer, so that was another set of difficult adjustments eliminated. We replaced the keyboard logic and LED scanner with a microcontroller which made remote control of the instrument possible by 'reading' the LED display and 'pressing buttons' on the fly. This meant the instrument could be tested automatically in production although it didn't have an external HP-IB interface.*"

One area that transferred directly from the old 4935 was the single-chip microprocessor. The firmware running on this was stored in a mask-programmed ROM, so no changes were possible. This did have the advantage of saving on development time and cost, but presented some constraints. Andy Batham and the team found some ingenious workarounds[16]. Sticking with the old processor also ensured the new instrument had a very similar personality to its predecessor.

The new instrument had a more compact case and a stronger metal handle. The membrane keyboard and integral LED annunciators gave a cleaner more ergonomic front panel. Furthermore, the improved design cut over $1k off the list price at launch. It cost $2.9k including the rechargeable battery pack, compared to $4k for the 4935.

The Q-TAP delivered, and the 4934A hit the market at the end of 1988. Customers loved what Andrew Wilson, the Product Manager, described as the "*Harley-Davidson of TIMS boxes – old wine in a new bottle.*" The new TIMS rather confounded its various critics: the digital zealots in R&D who thought it was passé, the hired consultants[17] who advised HP against low-cost telecom test products for the North American market, and the business gurus in Division marketing who thought it would be a poor return on investment. The 4934A turned out to be the highest volume instrument ever designed and built at South Queensferry! Andy Batham and the team celebrated the end of the project by having a day out climbing Ben Lawers in the Scottish Highlands. Apparently a 4934 went too, being portable.

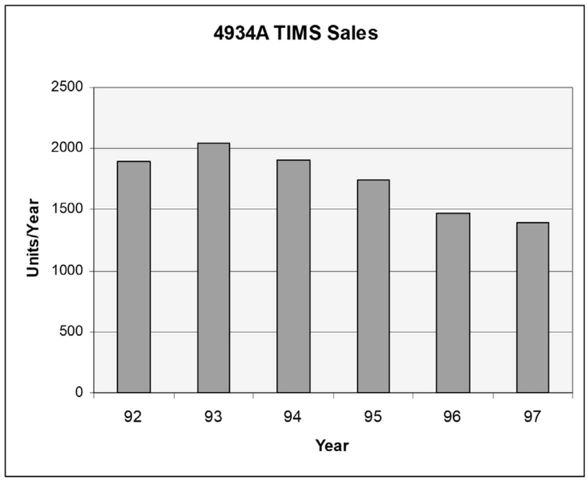

Full sales figures for the instrument no longer exist, but we do have records for the six years between 1992 and 1997 that give an idea of its sales. In December 1996, the manufacturing team had an informal lunch to celebrate 14,000 units shipped since its launch in 1988. The celebration also marked the start of volume production of a 200kHz/RS-232 version.

During its 10+ years of production life from 1988 to 1999, it is likely the Division shipped well over 16,000 units worth more than $45M in revenue. Despite its success, Division management were reluctant to authorise a CCITT version to replace the old 4936A, so the product team engineered an "under the bench" project to create a special option, the 4934A Opt J01/2 which incorporated the psophometric filter and Siemens 3-pin connectors. This turned out to be useful when the 4936A was discontinued in 1992.

[16] Stuart Connelly, the Product Support Engineer, recalls one in particular. The 4935 had two exclusive measurement options, PA/R or noise-with-tone. These were programmed by a link which the processor read on powering up. As the 4934 had both measurements, when the user pressed the relevant button, it was latched and the processor put through a full reset so it woke up again thinking it was the other 4935 option. It worked, but was unconventional to say the least!

[17] CORRAO & Associates Report c.1988

The 4934A TIMS was the last analogue communications product designed at South Queensferry and it was a glorious finale. The unit preserved in the National Museum of Scotland (NMS T.2010.71) lacks a serial number so its date of manufacture is currently unknown.

Retrospective

HP's involvement in this market over nearly 40 years is intriguing. Initially exploiting the Company's expertise in audio signal generators and AC voltmeters, by the 1980s and 90s the business had become a commodity market driven largely by price. HP was the world's premium test equipment manufacturer, famed for its high-end microwave and high-speed digital instruments: markets with very high barriers to entry and products that commanded top prices. In contrast, anybody could make a low-end TIMS, and tellingly, all the competitors for the 4934A were small companies such as Ameritec, Navtel, Northeast and CXR, not exactly famous names in the test equipment business. Nevertheless, HP's offerings were competitive, and it is thought the company had more than 40% share of the market, a dominant position. It demonstrated the power of a quality brand. It was safe to buy from HP and was probably a bit like the computer business where it was said that nobody ever got fired for buying IBM.

Acknowledgements

The following former employees kindly assisted with this chapter: Finlay Mackenzie, Robin Myles, David Guest, Stuart Connelly, Alex Ballantyne, Andrew Wilson, Gordon Rhind, David Grieve, Mike Kerr and Andy Batham.

7

Chapter Seven

The RATES System

Operational Equipment is the term used in the telecoms business to describe any of the facilities that carry or control the revenue-earning customer traffic. The equipment and software have to be very reliable while conforming to a whole range of standards, sometimes referred to as the Equipment Practice, which defines the physical dimensions and rack-mount hardware, electrical connections and power supply, software operating system, environmental specifications and even the colour of the front panel! The operational equipment was specified to have a working life of 25 years in earlier days, though that eventually became a bit ridiculous given the rate of technical change. Walking into a telephone exchange in the 1970s or 1980s, you were immediately struck by the uniformity of everything. In British Telecom, the equipment racks were all painted in a pale yellow called Light Straw[1], and the cables were a pale cream colour. Aside from this, test equipment such as that provided by HP at South Queensferry, usually retained its standard commercial appearance even if it was an approved model with, for example, a specific tester number. Test equipment didn't need to conform to the equipment practice as it didn't carry customer traffic.

Because of these strict conformance issues and "type approvals", monopoly telephone companies tended to buy all their operational equipment from the same few suppliers who also undertook new equipment development in close association with the service provider. In the UK, the British Post Office (later British Telecom) bought mainly from companies like Plessey, GEC and STC, while in the Bell System in North America, the majority of equipment was supplied by AT&T's manufacturing subsidiary, Western Electric (WECO) and Northern Telecom. This presented a considerable barrier to entry for any other equipment supplier, until the advent of deregulation and competition in the telecom market from the mid-1980s onwards.

[1] Apparently BT switched to Light Straw from earlier battleship grey as they found the lighter colour saved about 25% on lighting costs.

Despite these obstacles, South Queensferry decided around 1980 to venture into what was in effect the operational equipment market, with a new development called Remote Access and Test or RATES. All pretty much uncharted territory for the Hewlett-Packard Corporation, which tended to avoid large contracts involving custom engineering in favour of data-sheeted products. No doubt this new Scottish venture was viewed with some suspicion at Corporate headquarters, but by then the Division had acquired the reputation for "doing unusual things"!

What did RATES do?

In Chapter 6 we explored South Queensferry's family of Telephone Line Analyzers and how they were used to measure the various transmission parameters of an end-to-end telephone connection. The principal requirement was to measure the line to ensure it would be suitable for a high-speed voice-data modem. Before the wide availability of direct digital and broadband connections, this was the main way of establishing data connections between computers and terminals. As the transmission had to be sent as a modulated signal through the restricted 3.1 kHz bandwidth of the telephone channel, the performance, in terms of signal-to-noise ratio and pass-band characteristics, was quite demanding to minimise errors on the data signal.

To meet these demanding specifications, the telephone company provided a leased-line or private circuit between the two customer sites and could guarantee that the performance of this "nailed-up" circuit would meet the line specification since it bypassed the normal telephone switching equipment. These specifications were sometimes referred to as the "tariffed parameters", since the service provider charged a premium for these circuits compared to normal switched telephone service.

To allow these special circuits to be reconfigured and accessed for testing, they were terminated at each exchange on a "jack-field" or "cross-connect panel" so that various circuits could be interconnected using "jack-cords". Testing could either involve breaking into the circuit, or simply monitoring the signals on the circuit by bridging it. The measurement connections were then routed via the jack panel back to the test desk in the exchange, equipped with an audio oscillator, level measurement and other basic electrical test. The testing process was completely manual and often involved communicating with staff at other exchanges to make end-to-end measurements. These photos show the Access and Test facilities in a large exchange in Edinburgh in the early 1980s.

Jack-field and Cross-Connect

Record Cards

Test Desks

As the market for "private circuits" expanded in the 1970s and 1980s, major cities would have thousands of leased lines terminating on jack-fields in the various exchanges. In the UK alone, it was thought there were around 200,000 private circuits in use by the early 1980s, and probably double this later in the decade. Keeping track of all the interconnections and which customer's circuit was being carried on which connection, required a massive database of information. Again, this had been traditionally a manual process using files of Record Cards that had details on the customer circuit. Before any testing could be done, the circuit had to be identified using the record cards.

Clearly, the whole process was ripe for computer automation both of the record keeping and accessing the lines to be tested. There would be an enormous improvement in efficiency, down-time and responsiveness to customers. Over in North America, a system doing exactly that had been rolled out during the 1970s. It was called Switched Access Remote Test System (SARTS) supplied by Western Electric (WECO). Since the customer traffic had to be routed through the access switching (which replaced the jack-fields and jumpers), SARTS was treated as operational equipment. The system also had the power to intercept the traffic.

The size of the potential market became apparent as Finlay Mackenzie (former South Queensferry General Manager) recalled: "*An HP sales engineer in Santa Clara who had previously worked on installing SARTS told us that around $1B had been invested in the WECO system in the USA during the 1970s and early 80s.*" It was certainly a very big market, but given the operational equipment barriers, how much of it was accessible?

A New Venture

In the late 1970s and early 80s, the business climate was not good, particularly in the UK, and the Division was looking for some new opportunities in order to grow the operation. A year or two before, South Queensferry had entered the computer controlled remote monitoring business with its early FDM surveillance systems (as described in Chapter 5 on the SLMS family). The FDM surveillance system software sat on a software platform called Telecom System Supervisor (TSS-1000). This buffered all the Division's telecom application software packages from the HP1000 computer's RTE (Real Time Executive) operating system. This had two advantages: any changes to the RTE software only required appropriate changes in TSS and not in every application package; it also allowed any or all of the Division's software packages to co-reside on the same HP1000 computer system. This meant that one computer could run the FDM software as well as all of the future RATES software packages. So the Division already had an important ingredient for a SARTS type product.

Also during the 1970s, the Division had success with the 3770A/B Telephone Line Analyzer, particularly with sales to the British Post Office, and hence understood the line measurements required in this type of system. In addition, the extensive portfolio of TIMS instruments in the HP catalogue (see Chapter 6 on Telephone Line Analyzers) provided plenty of design knowledge within the Company.

The major challenge was the access switching itself, since nobody at the Division had experience of this and it was doubtful if anyone within the HP Company had any knowledge. Bearing in mind that the customer's circuit would be looped through the switching modules, they needed to have the reliability and maintainability of operational equipment. Some key requirements were:

- The relay switches in the access modules needed to provide a dependable low-loss connection even though the relay might not be operated for months or even years. This was a stringent requirement since relays often rely on the wiping-action during operation to maintain low-resistance at the contacts.
- The access switch needed to be fail-safe so that if power to the module were lost, the customer connection would be unaffected.
- The module needed to have make-before-break connectors, so that if an access switch card in the module were pulled out, continuity would be maintained on the customers' circuits.
- The architecture of the access switch needed to be designed so that it could be expanded and adapted for different exchanges and provide an acceptable level of "blocking"[2] for circuit access while minimising the number of expensive high-reliability relays used.

Finally, another challenge for South Queensferry was convincing the customer that the Division understood these requirements and could deliver them, without a track-record in supplying operational equipment.

In early 1979, the Division heard that a major US-based company, ADC Telecommunications[3] near Minneapolis, was interested in selling its access and test system, which presumably had been designed to compete with the SARTS system. South Queensferry needed some reliable access switching so it seemed like a good opportunity and a visit to Minneapolis was organised. ADC emphasised the need to use the right relays in the switches for reliability. The product looked promising and Finlay Mackenzie recalled that the then R&D Manager, Bob Coackley, was very keen to do a deal with ADC. Negotiations took place between ADC and HP legal in Palo Alto, but the sticking point was that ADC wanted HP to take responsibility for all their pre-installed systems. A further meeting, this time in the agreeable surroundings of the Bahamas, resulted in a deal on HP's terms.

The deal, signed in mid-1979, gave South Queensferry the rights to the design and know-how of the access switches, and involved an up-front lump-sum of around $400k and quarterly royalty payments based on the number of switch and measurement cards shipped. Initially, the Division bought some switching cards, and sample hardware of the Test and Measurement Unit (TMU).

However, after further investigations on the equipment bought-in from ADC, the R&D team decided to design their own system. Peter Locke, a software designer, commented that the software and microprocessor control had too many "bugs" and faults to give

[2] Blocking is a feature of any efficient hierarchical switching configuration, and occurs when the setting-up of one access path might prevent the setting-up of another path since it would require switches already in use for the first path. A completely non-blocking switch is possible but very expensive as so much hardware needs to be duplicated.

[3] In the early 1970s, ADC supplied prewired, jackfields, wired assemblies and test equipment for telephone operating companies. By 1976, ADC had become the largest independent supplier of test boards for telephone systems in the United States.

repeatable results. "*Some of us felt the product hadn't been evaluated properly before we made a hasty decision to buy it.*" Meanwhile, two other engineers, Robert Duncan and David Guest, evaluated the hardware. Robert recalled, "*The physical architecture was ghastly, and was cabled at the front of the racks which was unacceptable to customers.*"

Despite these shortcomings, the ADC system did provide the Queensferry design team with the basic architecture for the access switching, meaning work could start with less product investigation. Robin Myles (Product Marketing Manager) commented,

"*Around this time, David Dack and I went on a trip to the USA and realised that the new system needed to be compatible with SARTS as so much was already installed. We also realised that the added costs and royalties associated with the ADC design would make us uncompetitive on price.*"

So work started on the development of the new South Queensferry system.

A British Post Office Contract

Sometime in early 1980 a meeting took place in London with the British Post Office to discuss their plans for automating access and test. This was initiated by Dennis Tomkins, the Major Account Manager in the UK sales company dealing with the Post Office. The sales force had become aware that a tender would in due course be issued for a major contract to supply access and test equipment and that the preferred supplier was likely to be Plessey, a company already well acquainted with operational equipment. South Queensferry's embryonic offering looked a very distant second choice at that stage.

Back at the factory worked progressed on designing some access switching cards using the relays recommended by ADC. As Robin Myles recalled, this wasn't going to be good enough:

"*The Post Office told us we had to use specific relays with contacts made of almost solid gold, or they wouldn't give us any business. At the same time, the UK field sales told us we were struggling against the Plessey offering, and that a $6M contract was likely to be awarded in 1982.*"

Then followed a lucky break, as Finlay Mackenzie remembered:

"*HP UK had just completed a new building in the Reading area, and as the new General Manager at Queensferry, I'd been invited to the official opening. One of the other guests was a senior executive from British Telecom. I sat next to him at lunch and he was a very good contact. When the RATES deal appeared to be slipping away, I phoned him and suggested BT should view what we had to offer, as it had a lot of potential. We believed that success with BT would help us win export orders. He agreed to come to South Queensferry, and during a lunch at the Champany Inn, we told him about our overall systems capability. Afterwards, back at the factory David Dack, Lawrence Lowe and Peter Locke did some presentations on our RATES solution.*"

This intervention had a result, as shortly after Finlay heard from the UK sales staff, that BT would like to meet again and would postpone the closing date to allow HP to bid, although

Plessey remained the preferred bidder. Very soon a 12" high pile of paper arrived at South Queensferry, the tender documents from BT, and Graham Madeley and Reid Urquhart had a short time to respond.

This turned out to be a challenging process as the tender had to include the technical and delivery responses from South Queensferry for the switching and measurement hardware and also, far from simple, the definition of the software we would deliver to control the system. Meanwhile, the UK sales management, who conducted the negotiations with BT, had to decide on the commercial side, had to work with their colleagues from computer sales as this formed part of the package, and last but not least, decide how the system would be installed and supported, and by whom. It was probably the first time HP UK had undertaken such a complex bid and almost certainly the first time the HP Corporation had tendered for telecom operational equipment. It was a marked departure from HP's usual sales process for selling instruments, usually referred to in the factory as the "box business".

Finlay commented that the UK Test and Measurement Sales Manager, Roger Thornburn, really understood the needs of the contract and put the resources and people in place to support the process. The tender was delivered to BT on the 16th March 1983, and Finlay described the whole exercise as "*a magnificent example of teamwork*" between the product development team at South Queensferry and the sales staff at Winnersh and Redhill. Graham Madeley, from product marketing who had the main responsibility for managing the contract at the factory, recalled the considerable amount of work required to complete the proposal:

"*The Division did all the work on the technical configuration for each of the BT exchanges, and the complete equipment list for each site with prices. I remember one year making over 60 flights to London as we worked on the detailed requirements of the contracts.*"

The London-area sales office staff were also frequent visitors to the factory and became familiar faces around the place.

All this work paid off, and on the unlikely date of Friday 13th May, the team heard that they had been awarded a major part of the first BT order for RATES. It was a significant achievement, considering HP UK had not previously bid for this type of contract. The access switching part of the contract was split between HP and Plessey[4], but the computer, software development and the test and measurement hardware all went to HP. This amounted to £3M ($4.5M) out of a total contract of £4M ($6M).

The jubilation was tangible. At many of the larger UK sales offices, the win was announced over the "Tannoy" systems that Friday morning. Finlay Mackenzie's satisfaction was recorded in the July 1983 "*Readout*" newsletter:

"*As you can imagine, I was delighted to hear we have won the RATES order, which is four times larger than any other single order ever received here at Queensferry. I can assure you that our 'easy-to-use' software combined with our network ability played a very large*

[4] The Plessey offering was based on a North American design (TTI SAS). The access switching was based on a 6-wire grab which suited North American systems whereas BT's specification requested a multiple of 4-wire grabs for most efficient use of the access switching. The HP system used an 8-wire grab so that a 4-wire circuit could be split and terminated in both directions.

part in winning the order, particularly with our unique combination of measurement and computer systems. Also, in talking with senior staff at BT, we learned that they were particularly impressed by HP. They commented favourably on our high-level of commitment, our team spirit, and our motivation. In fact, this infectious enthusiasm of ours should never be underrated when it comes to the secret of our success. The greyhound overtook the tortoise at the last minute!"

Roger Thornburn from the sales side made a similar comment:

"*Everyone became infected with the spirit of the project as our work progressed, and our perseverance right up to the last minute really paid off – even though there were times when our chances looked pretty slim.*"

Then the hard work began. Deliveries of the equipment were due to start in November 1983, so the product development needed to be completed quickly. By then, the Division had a large team working on the project – 21 engineers in total on design, software, marketing and support. Graham Madeley commented:

"*We were initially a bit naïve about what was required to execute such a large contract in terms of documentation and logistics. I think the Product Support Manager, Hugh Smith, was one of the few who understood the magnitude of what we had got into. As time went on, it got easier as we developed systems and procedures.*"

This cartoon, circulating round the team at the time, summed it up!

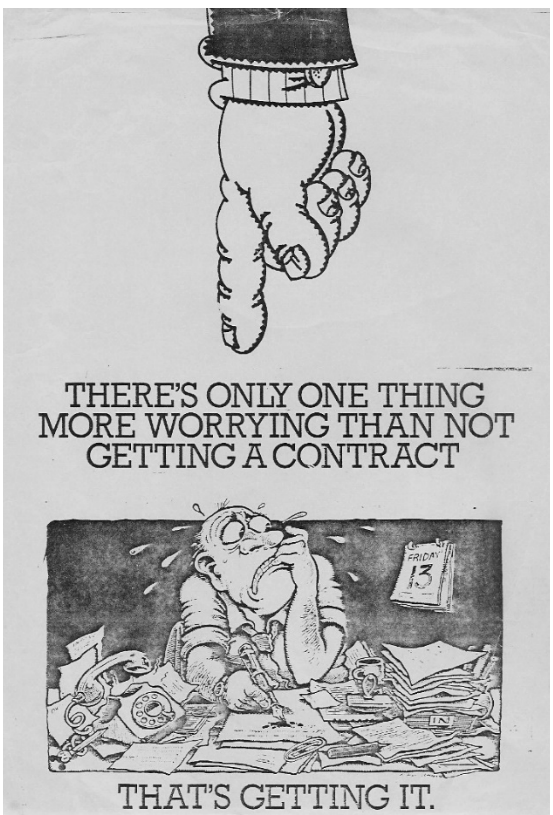

Derek Peek, a key project engineer and later Project Manager from the sales side, agreed:

"*We were naïve when running a contract of this type, and BT exploited this. I imagine Plessey would have had a project manager in overall command with the necessary authority to control all aspects of the contract. Our internal structures were all set up for designing, building and selling boxes (instruments), not systems.*

"On the positive side, we did have a lot of continuity on the team. Keith Mitchell was the BT Account Manager throughout the RATES project and did a great deal of work on the first and later contracts. He reported to Peter Kitson who was also responsible for the BT account throughout the contracts. The relationship between our team and the factory was great, and everyone pitched in to get the job done in the time-honoured HP way. If it was not for myself, Ken MacDougall, Derek Milne (at South Queensferry) and others willing to work around our internal systems and structures, things would not have gone as smoothly on the first contract."

The System

At this point it would be useful to look at the anatomy of RATES. Although this chapter has so far used the name RATES for South Queensferry's system, it was actually the acronym used by British Telecom for the system they rolled out nationally during the 1980s. The codename for the HP product was RTS 10 (Remote Test System No.10) and the product number was 37100S. All the 371XX numbers were reserved for the modules and cards used in RTS 10. Because the BT contract was so significant, over the years the system was simply referred to as RATES at the factory, and the RTS 10 title lapsed.

The RATES hardware consisted of three main elements:

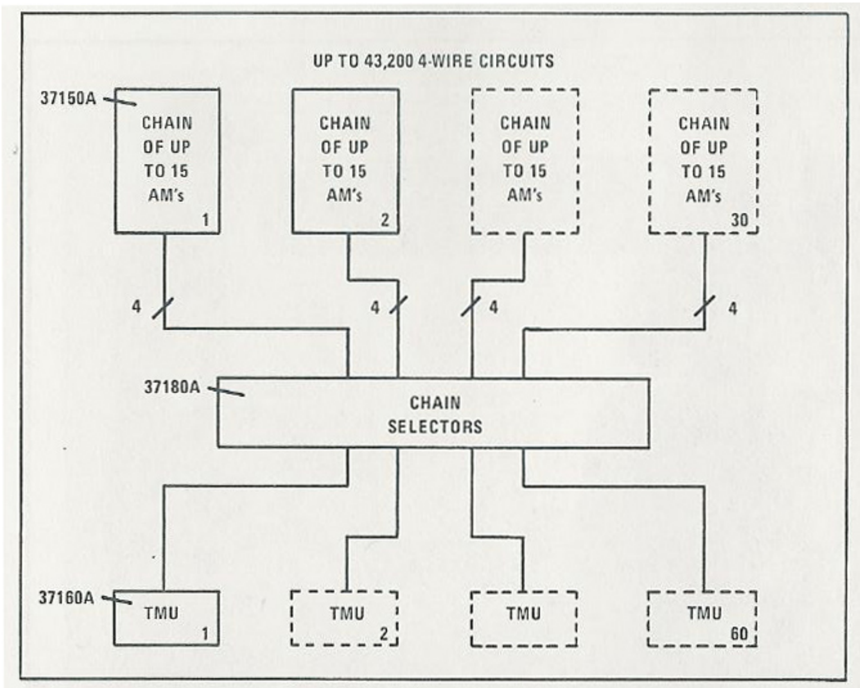

The key element was the 37150A Access Module, a 19" rack unit that contained several switching cards with the all-important gold-contact relays. The customers' private circuits were wired through these modules via multi-way sockets on the back panel. Each Access Module could handle 96 4-wire circuits through its four access cards. A chain of 15 Access Modules could be connected together to provide a total capacity of 1440 4-wire circuits. For smaller systems, this chain of Access Modules could be connected to up to four 37160A Test and Measurement Units (TMUs). The TMUs were also 19" rack units containing a range of cards for different measurements including, transmit/receive level measurements, noise, multi-meter (AC/DC voltage, resistance etc.) and signalling.

In common with other "operational equipment" connected to long external telephone circuits, the RATES hardware had to operate correctly in the presence of strong unwanted signals (longitudinal or common-mode signals), and withstand transient surges of several hundred volts without damage. Interestingly, the selective and weighted noise measurements used the new Queensferry-designed digital filter chip, first used in the 3776A/B PCM Test Set, as described in Chapter 8. The TMU was also the focal point for control of the access/measurement subsystem, and provided a port for connecting a local user terminal which interfaced back to the central control computer.

On larger systems like the one shown in the diagram above, several chains of Access Modules could work through 37180A Chain Selectors to interface to a bank of TMUs. With this arrangement, very large exchanges could be accessed.

Around 1985, another couple of variants were introduced which used a wider 23" rack module. One of these was the 37140A which combined a 96-circuit Access Module with a TMU, so that a complete subsystem for a small exchange could be provided in a single module with less interconnection and lower cost per access.

Another variant was the 37130A Dual Access Module giving a capacity of 192 circuits. Again it reduced the cost per access, but with slightly poorer blocking. The original 37150A Access Module had been designed with an option slot next to each switching card for accessories such as a Jack Access Module. Very few of these accessories were sold so the space reserved for them was redundant[5]. This over-design was eliminated in the 37130A/140A card cages

All this equipment was designed to run off the -48V DC supply in the telephone exchange. A whole range of circuit cards could be configured in these modules to satisfy a particular customer requirement. A full list of the various cards is given in the Product Catalogue in Appendix 1.

A major part of the RTS10 RATES system was the software running on the HP 1000 A-series computer. As mentioned earlier, this ran South Queensferry's Telecom System Supervisor (TSS-1000) which was built on top of HP's Real Time Executive (RTE) operating system. Several application packages for RATES were developed by the software design engineers. These included the main control software, circuit record card software, test point assignment and auto-routining software.

Robin Myles recalled that the software development was initially a bit fraught as it had to be completed quickly and the original tender specification wasn't precise enough:

"*BT were tolerant of the problems as they realised they should have provided a better definition. The man-machine interface wasn't what they expected. However, we found it was better to let them evaluate a prototype and work from there, which was more efficient than working from documents. BT engineers were regular visitors at Queensferry, working with the software guys to fine tune the operation.*"

The HP 1000 computer provided centralized remote access and test, with the terminals connected directly to the computer, or via a TMU at a remote site to provide the same power as the centrally connected terminal. This meant any circuit could be accessed from any terminal, and the circuit record database could be modified from any site as the changes were made.

[5] Although the accessory slots were not used, the customer circuit still had to be routed through a second set of make-before-break connectors, inevitably compromising reliability to some extent.

Quite a lot of effort went into making the system easy to use, by means of screen graphics and soft-keys. The technician was asked to make a choice from the eight soft-keys along the bottom of the terminal screen, as shown in the photo below.

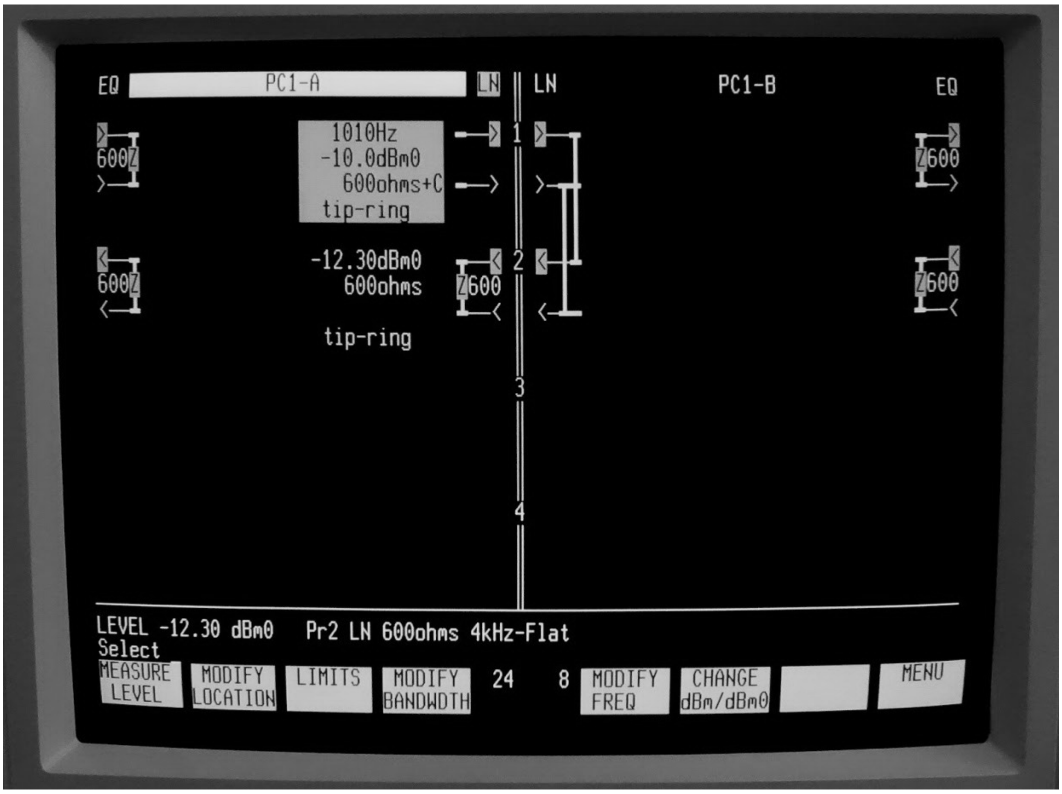

On pressing a soft-key, the labels on the screen would change and new set of eight soft-keys were then presented to make another choice. Thus, the technician was guided through the test sequence and only needed to enter information about the specific circuit or test point. The soft-keys also provided an elegant way of adapting the user-interface to the particular system configuration in say the TMU.

More BT Contracts

BT was clearly satisfied with the first consignment of RATES hardware and software which they started to receive at the beginning of 1984. In October 1984, South Queensferry was awarded another large contract worth a further £3M ($4.5M). BT continued to use both HP and Plessey to supply equipment, no doubt a prudent strategy to avoid becoming dependent on one supplier for this important system.

A press announcement in November 1984 noted that the first phase of RATES had covered 63 exchanges and this second phase would further extend the number, with the long term objective of encompassing 2300 large and small exchanges by 1987. The second phase increased the number of installed computer sites to ten.

By August 1986, further orders from BT meant the total business had risen to over £15M ($23M) after discounts, representing over 30 computers and 700 exchanges covered by the RATES system.

HP 37150A Access Modules in Card Cage

HP 37160D Small-system Access/TMU Combo

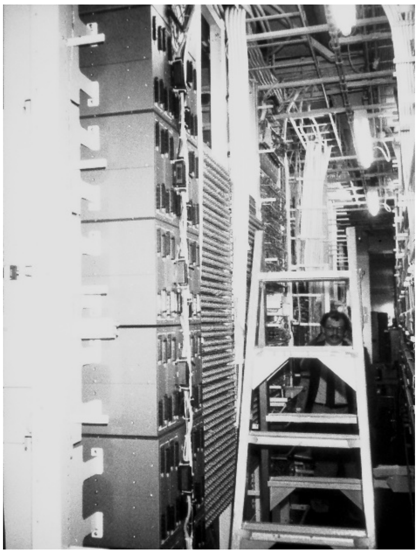

Large RATES installation in BT

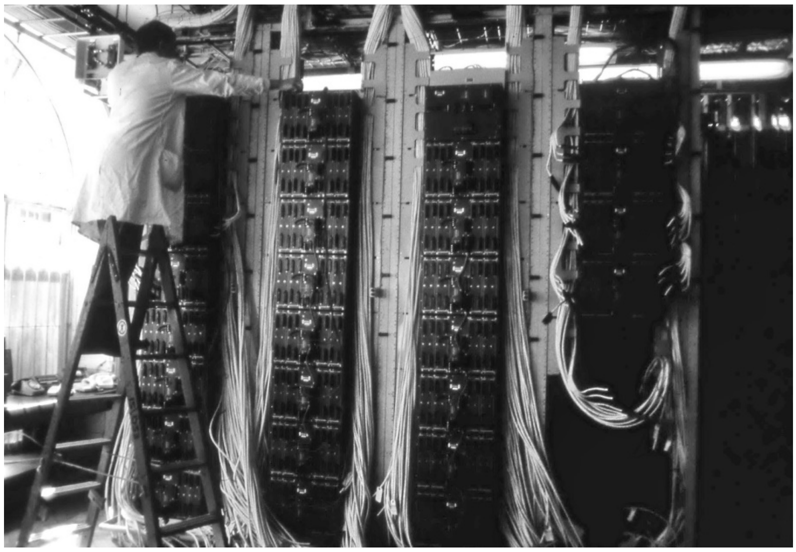

Installing the Cabling

RATES installation nearly complete

Simulated RATES Operations Centre

The *Readout* newsletter of October 1986 recorded the celebration following the latest order:

"*The sales team was in euphoric mood at the champagne celebration at HP's Redhill office near London. Most of the team had worked together on the BT business since 1983 – in itself an unusual length of time for a project team to stay together. As Project Manager, Derek Peek pointed out, 'Everybody pitched in on this business – we've all spent some long evenings on this job, getting to know each other's eccentricities and tastes in Indian food!'*

"*But such an unusually large project demands unusual commitment, and as Stuart Ritchie, Computer Sales Engineer on the BT account, said, "It all goes to prove that project business needs teamwork. If we hadn't been working together all this time between the factory and the field, and really getting to grips with BT's requirements, our task would have been much harder.*

"*The end result of all this is that BT engineers are delighted with the RATES system and say it is one of the most successful projects they have ever seen. An important reason for this high customer satisfaction level, says Ken Stead, Installation Manager, is the fact that HP decided to set up a special team to look after the complete system.*"

Several of the 700 exchanges mentioned above had very large installations of RATES equipment testing thousands of lines – racks and racks of access equipment stretching up to the ceiling in the exchange room with thick bundles of cream cables entering from the top. All the HP equipment was painted in "light straw" to match BT's operating equipment rather than HP's colour schemes! Archive pictures of these installations gives an idea of the logistics and installation work required to complete these contracts successfully.

Derek Peek was responsible for the installation, commissioning and support of the system. He had done his apprenticeship with EMI in the installation and maintenance division, and before HP had worked for London Transport installing communications for the then new Victoria Underground Line.

"*I had a team of six first-class engineers who worked in teams of two at each site. Ray Longman took responsibility for installing the HP 1000 computers and the software. We found the attitude to RATES in BT varied across the country. Some areas such as Glasgow embraced it and the guy there used it from the start and became very knowledgeable on the system, whereas the guy in Norwich switched off the HP 1000 after it was installed and only switched it on if we came to do software upgrades! I think the resistance may have been due to the possible job losses the system would cause, but in the 1980s there was also a fear of computer technology and the guys testing private lines were happier with manual testing.*"

In February 1989, BT issued a final invitation to tender for further extensions to the existing systems to accommodate growth, and also a new project to equip up to 1300 small exchanges around the UK with RATES equipment. The earlier systems were referred to as "RATES A & B", while the new small exchange application was described as "RATES C". This final tender was called Phase 4.

A copy of HP's tender response for Phase 4 of RATES, and the BT tender specification, has survived. It is contained in three smart ring binders in dark blue, packaged in presentation boxes. HP's response is dated April and May 1989, and gives an idea of the work required by the field and factory to compile this kind of submission, and the commercial decisions necessary to present a competitive bid.

BT provided details of the 1300 "small exchanges" and the number of 2-wire circuits in each. This showed that nearly half had 192 circuits (96 4-wire circuits) or less, with the greatest number having only 96 2-wire circuits. Some of the small sites included familiar local places like South Queensferry, Linlithgow, Alloa, Dunfermline and Lochgelly.

With such small numbers of accesses, the previous RATES architecture of a separate TMU module and Access Module would have resulted in a high cost per access. In the standard TMU (37160A/B), each measurement function was allotted a separate plug-in card so that it was easy to configure the measurements to exactly those required by the customer. For RATES C, the engineering team at South Queensferry took the standard TMU and redesigned the measurement hardware to combine two measurements per card. This reduced the number of TMU cards from nine to five, freeing four slots for the 37151A Access Cards. This meant a single 19" module could now handle all the measurements needed and provide access to 192 2-wire circuits. In his cover letter in response to the tender, Peter Kitson, the District Sales Manager, claimed this would make the solution nearly 70% cheaper than the previous RATES B configuration for the smallest exchanges. These modified modules were all networked back to the main RATES computer centres already installed, and HP was keen to emphasise that this new RATES C hardware would be fully compatible and have all the functionality of the earlier installations despite its smaller size.

To justify the re-engineering and the heavy discounts of 40 to 50% on list price, HP was bidding for a large share of the business on offer: "*All the RATES C business for £4.5M ($7.2M) and a minimum of £7.8M ($12.5M) of the Phase 4 growth business, net of discounts*". The Phase 4 Growth contracts were large extensions of existing systems, in particular for Moorgate and Monument Exchanges in the City of London for which HP quoted about £3.5M ($5.6M) for an additional 116,500 2-wire access points (employing 2430 x 37151A Access cards). Overall HP quoted a total of nearly £13M ($21M) for the Phase 4 Growth Project, with a net value of £7.8M ($12.5M) after a 40% discount. BT continued to place some of the access switching and installation work with GEC Plessey Telecommunications (GPT), and in general where HP or GPT access equipment was already installed, extensions would go to the same company. In the above tender, HP did not quote for exchanges with GPT equipment installed. HP did have the upper hand however, since they had supplied all the control computer equipment and software, so bidders had to conform to HP's interfaces and communication protocols[6], and equipment had to be previously type-approved by BT.

The final BT contracts were awarded in summer 1989. The equipment was delivered during the first half of 1990, and marked the conclusion of the BT RATES installation apart from some minor additions and modifications, particularly to the control software. In a lab report from January 1990, Mike Hurst, software project manager, noted that BT engineers had visited in December to discuss further enhancements and modifications to

[6] This was the N3 interface providing the communication between the HP 1000 computer and the remote Test Access Equipment (TAE) usually over a modem link.

the software, and around that time the Rev. 5 version of the RATES software was released. Work was also going on to test the multi-processor software for linking the various RATES computers. BT engineers were due to undertake an acceptance test that month. While the RATES hardware had been heavily discounted to win business versus Plessey, the software was much more lucrative for HP as the arrangement was non-competitive, and the business continued for several years after the last major BT contract.

The value of the BT RATES contracts was probably over $50M including equipment, installation, training and follow-up software enhancements, representing a major chunk of business for the factory and the sales team in the South East. Phil Gibbins, the UK Account Manager for BT summed it up quite well in an article in the August 1989 issue of *Readout*: "*During the life of the RATES Project, the factory and field have built an excellent working relationship. I believe it's this that gives us a competitive edge.*"

For many of those who worked on the RATES contracts in the factory and the field, it was a highlight in their HP careers. As Derek Peek recalled, "*I thoroughly enjoyed working on RATES and with South Queensferry, and it was my best time in Hewlett-Packard.*"

On the back of the BT business, there was a further contract with the UK's only independent locally-operated telephone company at the time, Kingston-upon-Hull Telephones, later Kingston Communications. Formed in 1902, it is known for its cream coloured telephone boxes in the city and was one of the earliest to provide mobile and broadband communications in the UK. The RATES system ordered by Kingston was broadly similar to the BT equipment and amounted to a contract of $0.5M to $1M, proportional to its size. The initial contract was for 10 sites, followed on with a further 10 sites or more. The most memorable part of this deal was the sales pitch, as Graham Madeley remembers:

"*The presentation was given by Tom White who headed the UK telecom specialist sales team in the mid 1980s. It was one of the best pitches I saw on RATES and he brought out all the key selling points very effectively.*"

This was significant, as a few years later Tom White was invited to become the Marketing Manager in South Queensferry where he was a strong advocate of the systems business in the "systems versus boxes" debate. He set up the Telecom Systems Division at Queensferry in 1994.

The Rest of the World

When the first BT RATES order was won in the summer of 1983, Finlay Mackenzie wrote a very upbeat article in the *Readout* newsletter, titled "*Giant Order Points the Way*".

"*I would like to share with readers my thoughts on the significance this order has for our future development at Queensferry and the lessons to be learned from how we won it. Over the last three years I have stated many times the telecommunications test business of the future will become a SYSTEMS business, and this is part of our strategy to design, sell and support systems solutions. Winning the RATES contract shows that we have been preparing not for some distant horizon, but for the contracts that are being awarded in 1983. As far as Systems installation is concerned, the future is now!*"

"The combination of HP computers and HP measurement systems is a capability which our competitors find hard to match – and impossible to build up overnight. RATES is a motivational milestone on the path we have taken to provide systems solutions for our customers. Nothing breeds success like success and we shall continue to expand our development efforts in this vital area. Winning this order on our home territory is bound to help our chances of winning export orders for RTS 10 – and imagine how difficult it would have been selling our system in world markets if our own national organisation had not felt confident enough to invest in it. I believe we will do well in North America with this system."

Having the BT contracts and an expanding scale of installation in the UK exchanges would certainly have been a considerable advantage. BT staff seemed genuinely satisfied with RATES and would probably have been willing to provide testimonials and host visits by other telecom operators to view the working system.

Despite this and the strong motivation in the factory team, the Division discovered fairly quickly it was difficult to win any further large contracts in other countries.

There was clearly nothing wrong with the technical quality of the system, rather the problems boiled down to two almost intractable issues. Firstly there was a reluctance by large telecom operators to place a substantial contract for operational equipment with a vendor for whom they had no previous track record. Secondly, and in a way related to the first problem, the local sales organisation often couldn't come up with a credible plan to handle the contract and support it, simply because they didn't usually do business that way. Furthermore, sales staff may not have had regular contact with senior management in the customer account, or even known who they were, as more often they dealt with engineers. However, the senior management were the key decision makers in a multi-million dollar RATES type system. Few could show the commitment of the UK sales team under Roger Thornburn, and this would count against HP in trying to get business. By comparison, the operational equipment manufacturers were very strong in this area.

As mentioned earlier, the North American market for these systems was enormous. Surely Queensferry could win some share of this business. In November 1984, Robin Myles, the Project Manager, and Ray Scott, a software designer, had a meeting with Bell Labs in Holmdel to discuss compatibility between RTS 10 and SARTS. In February 1985, a three-day visit to Bell South took place, this time Robin Myles was accompanied by David Dack, the Section Manager and Peter Locke who also worked on software. Apparently nothing materialised from these meetings, although it seems they were at the right level. Perhaps SARTS was already too entrenched in the Bell Operating Companies and there was a lack of interest from the field sales organisation.

To help promote the new system, the marketing group created a professional video. It was quite expensive, but not particularly effective, as Graham Madeley remembers:

"*We commissioned the services of William Woollard, a well-known British TV presenter of science and technology programmes, and had the video produced by the Heriot-Watt University TV studios. Woollard was polished and absorbed the script with astonishing speed. 'Here, in the shadow of the Forth Bridge, etc. …'. However, the video went down like a lead balloon in the US. It made us sound terribly British, and some customers asked for the video to be played at half speed so they could make out what was being said!*"

During the 1980s, South Queensferry established a North American sales support team in Colorado Springs to develop the Division's business. Headed by Tom Smith, a team of five or six engineers, some recruited locally and others seconded from Scotland, helped the field sales engineers demonstrate products and close deals. This closer coupling to the field meant there was a better prospect of winning a major RATES deal in the States. While the Bell business didn't look promising, during 1984 the field and factory identified a major opportunity with GTE[7], the largest independent telephone company in the USA.

The local HP sales engineer in Seattle, Jerry Ericson, had built up a very good relationship with one of the staff at GTE, Doug Campbell. Sales staff in other US regions, in particular Dallas (where GTE had its headquarters), Atlanta and Pennsylvania, joined in the project as its scale emerged and a field project manager was appointed. It was shaping up to be a repeat of RATES project in BT.

In February 1985, Robert Duncan, one of the RATES team working in North America, sent a memo to the Queensferry management:

"*I have just heard that HP has been awarded the GTE field trial in Seattle. This will now be the official field trial for the whole of GTE, not just the Northwest, and as such is excellent news. It represents a business opportunity of around $40M.*"

Robert Duncan recalled the close bond that developed between HP and the GTE staff:

"*I spent nine weeks installing the field trial equipment in Seattle and Oregon. I ate breakfast, lunch and dinner with those guys, and went camping with Doug Campbell and his family! The field trial was a success, but by that time there was a new manager at GTE, and she quickly figured out this was "end-of-life" stuff as everything was rapidly moving to digital. The whole project was scrapped and no one got any business.*"

Reflecting on this disappointing outcome, Robert commented that if the venture had been a couple of years earlier, then things might have been different.

"*On the whole, South Queensferry was five or even ten years too late into the RATES business and by the middle to late 1980s we were running into the digital era which signalled the end of these large analogue systems.*"

It was a fact that the huge investment in SARTS had happened in the late 1970s and early 80s, nearly ten years earlier.

Also in North America, there appeared to be some solid business in Canada, but as Finlay remembered, the field seemed to lack confidence to pursue the opportunity:

"*Alan Holdaway, the sales manager said to me, 'Finlay, you've got a tiger by the tail! How are you going to deal with all this stuff all the way across Canada from a factory in the UK?' I think a North American company, Halcyon or Hekimian, got the deal.*"

[7] GTE was the largest independent telephone company in the USA during the Bell System era and by the mid-1980s had acquired capability in long distance transmission and mobile networks. It was a big player in North America and the prospect of supplying RATES was a good opportunity.

Remembering a sales trip to Canada, Graham Madeley commented, "*It was a bit of a waste of time, as we probably weren't talking to the right people – not selling high enough.*" The only significant business in the USA was a small system for an operation called the Kentucky Emergency Warning System (KEWS) which provided critical communications links within the state. One reason this was successful, apart from the small scale, was that they also implemented South Queensferry's FDM surveillance on the same computer system.

It wasn't much better in Europe. The Division did some business with France Telecom, but none in Italy. Finlay recalled,

"*Spanish Telecom had a major demand for RATES, but the deal went to Telettra, a local operational equipment manufacturer, at the last minute. We also had some high-level meetings in Ireland with VPs of Irish PTT. At a lunch meeting, it was difficult to get discussions going. Maybe we were talking to the wrong people, or perhaps they just didn't want to buy from us. We didn't do any business there either.*"

It was a similar story in Scandinavia.

The South African Post Office (SAPO) selected the HP RATES system, and it was the largest installation outside the UK, no doubt helped by the close relationship Queensferry had over many years with the sales organisation in South Africa. The first contract was in 1983, and the customer was following closely the lead from BT. SAPO had a very distributed network with many remote sites where it could be unsafe to travel at that time, so the RATES sales proposition was effective. A total of three or four orders brought in revenue of over $2M. Robin Myles commented:

"*In my opinion one of the main reasons we did so well with SAPO over a number of years was the fantastic relationship one of the Queensferry support engineers, Derek Milne, built up with a key guy in South Africa.*"

As with each of the successful RATES deals, it was all about relationships between the factory, field and customers.

Finally, there was a successful sale to the telecom authority in Kuwait with sales of $1M-2M. During the first Gulf War in the early 1990s when Kuwait was invaded by Iraq, the RATES equipment was either taken or destroyed, so the Division got a repeat order once peace returned.

Unfulfilled Potential

With total sales of $80M to $100M, the RATES system will go down in the history of HP South Queensferry as a successful product, thanks largely to the business with BT. However, rarely again did the Division find that magic combination of teamwork, commitment and cooperation between field and factory, and the "infectious enthusiasm" that won the first BT contract against all the odds.

The 37100S was undoubtedly a superbly-engineered system, particularly the software, and it was one of the best solutions in the global market. However, even with the best mousetrap in the world, if you don't have the means to deliver it, you still won't do much business. And that was the problem: the RATES system was a misfit to the HP organisation and sales channel.

The Division seemed to be swimming against the tide, and it must have been a frustrating experience at times for the South Queensferry team. Perhaps there were easier ways to make money, a view that would not have been lost on the proponents of the "box business", of which there were a number at South Queensferry! Finlay Mackenzie recalls that after the demise of the GTE business in the USA, Dick Anderson, the Corporate Group Manager and his boss, contacted him and said,

"*The US field are wasting time on RATES. Would you take it off the US price list and out of the HP Catalogue?*"

If the Division had been able to replicate the UK field relationship elsewhere in the world, then perhaps the total RATES business could have been four or five times greater. On the other hand, perhaps that isn't realistic either. The UK relationship was intense, with key members of the sales team travelling to South Queensferry almost every fortnight to progress the contracts. It's hard to imagine that working internationally, except in the USA. Graham Madeley commented, "*If we'd had a dedicated sales force, then we could have cleaned-up.*" However, that might have been difficult given the entrenched protocols in the HP organisation – the field dealt with the customers not the factory, and there was a complex arrangement for assigning commission on sales.

There were essentially two different sales scenarios that faced HP. Firstly, there was the situation where the customer had already made a business case internally for a RATES type system and was in the market to find the right solution. An example was the BT contract. In this case HP needed to show it had the best package of hardware, software and support services, versus competitive offerings and the best value. Because of the scale of these contracts, the selling would still need to be at a fairly high level to gain credibility.

The second sales scenario was more difficult. In this case the telecommunications customer had first to be sold on the business case for such a system. This was not about product details, it was about understanding the customer's business and how the RATES system would improve their financial returns. This required special skills in high-level selling and building relationships, quite different from the business of selling instruments. This was where a "dedicated sales team" might have succeeded. Understandably, most of the RATES business was done by the first scenario.

There was also another factor limiting the opportunity in the late 1980s already mentioned in connection with GTE, namely the rapid migration to digital networks and services. The "window of opportunity" was closing for RATES and several other analogue communications products from South Queensferry. The heyday of the analogue private circuit, and the high-speed analogue data modems they supported, was the mid to late 1980s. After that, private networks started moving onto digital facilities, particularly in North America where the T1 leased line at 1.5 Mb/s became "de rigueur" for anybody building a private enterprise network. The T1 line carried both telephony and data, and the network equipment itself monitored quality without the need for a RATES type system. If an operator hadn't already installed access and test before the mid-1980s, then they were unlikely to, so it became an increasingly difficult sale.

There was a sort of transitional period in the mid-1980s when these old analogue leased lines were “trunked” through digital transmission facilities by using PCM encoding of the analogue channel. In the USA, a digital variant of SARTS was introduced, called DARTS (Digital Access and Remote Test System). Interestingly, South Queensferry won a contract in 1984/85 to supply several hundred 3776B PCM Terminal Test Sets to integrate into this system, however the Division never tried to upgrade its own RATES system in this way.

The 37100S was completely discontinued in 1994, as there was little prospect of new business, although there were further software enhancements for existing systems. On a happier note, some of the team who had worked on it transferred to a new venture, the HP acceSS7 monitoring system for Signalling System No.7. This led to a new division and a very successful business, as by then the HP infrastructure was more supportive of these contracts and it was an emerging market rather than one nearing its end.

Finlay Mackenzie’s prophecy that telecom test was becoming a systems business was proved at least partly true. Ironically, it became completely true in the twilight years at South Queensferry when systems became the only significant business left at the site.

Acknowledgements

The following former South Queensferry employees have contributed to this chapter: Finlay Mackenzie, Robin Myles, Robert Duncan, Peter Locke, Mike Hurst, Derek Peek, Gregan Crawford and Graham Madeley

8

Chapter Eight

The Family of PCM Testers

Pulse Code Modulation or PCM is the process of converting an analogue signal such as a telephone voice call into a digital bit stream. Although this general idea of analogue to digital conversion applies in many audio applications, for example CDs, the term PCM is very much associated with telephony, which was the earliest large-scale application of digitized audio.

Alec Reeves, working at STC (later ITT) in the 1930s, is usually credited with inventing the idea of digitizing the analogue signal and was granted a patent for PCM in the early 1940s. At the time, there was little commercial value as it was impractical to implement the idea with vacuum tubes or valves. That had to wait for the invention of the transistor. Much of the early work on transistors was done at Bell Labs in the USA, and it was here that the first development of PCM telecommunications took place in the 1950s. One of the engineers credited with this work was Barney Oliver, a prolific inventor at Bell Labs and later founder of HP Labs and close associate of Bill Hewlett and Dave Packard. In the US National Inventors Hall of Fame, he is cited as inventing PCM. This may not be strictly true, but Barney Oliver did probably invent the first "working" PCM system[1].

The earliest commercial use of PCM telephone systems was in the USA with the first installation in the early 1960s. By 1965, several thousand systems were installed. These were called T-carrier and involved PCM encoding a group of 24 analogue telephone channels and combining them by time division multiplexing into a 1.5 Mb/s digital bit stream which was typically transmitted across pairs of wires that had previously carried one analogue telephone channel. A similar T-carrier system operated in the reverse direction to provide two-way communication or, to use the telecommunications term, a 4-wire circuit. These early systems were mainly used as point-to-point connections between exchanges, simply to increase capacity.

[1] US Patent 2,801,281 held jointly with Shannon, filed 1946 and granted 1957 "Communication System Employing Pulse Code Modulation". A valve implementation with a large number of valves!

To cover longer distances, the signal needed to be boosted every mile or two by a digital regenerator which amplified and reconstructed the digital signal. If the regenerator received a very weak or noisy signal, it could create errors in the reconstructed bit stream. However, by designing the appropriate distance between regenerators, accumulating errors could be minimised. Even trans-continental and inter-continental distances could be covered without significant degradation in signal quality[2]. These first PCM systems therefore laid the foundations of everything we take for granted in today's global digital communications network.

Measurement Background

At the input of a PCM system, the analogue signal is sampled at roughly twice the maximum frequency contained in the analogue signal, and the sampled value then represented by a digital word of 8 bits or more. This process leads to a property called quantization. Whereas the analogue signal has an infinite number of possible signal levels, the digital representation can only take a finite number of levels determined by the number of bits used. So the 8-bit word can only have 256 possible discrete levels, and there will necessarily be an error between the quantized digital signal and the analogue signal it represents. This error can be made very small by using more bits – for example commercial audio CDs use 16 bits and very high quality may use 20 or 24 bits.

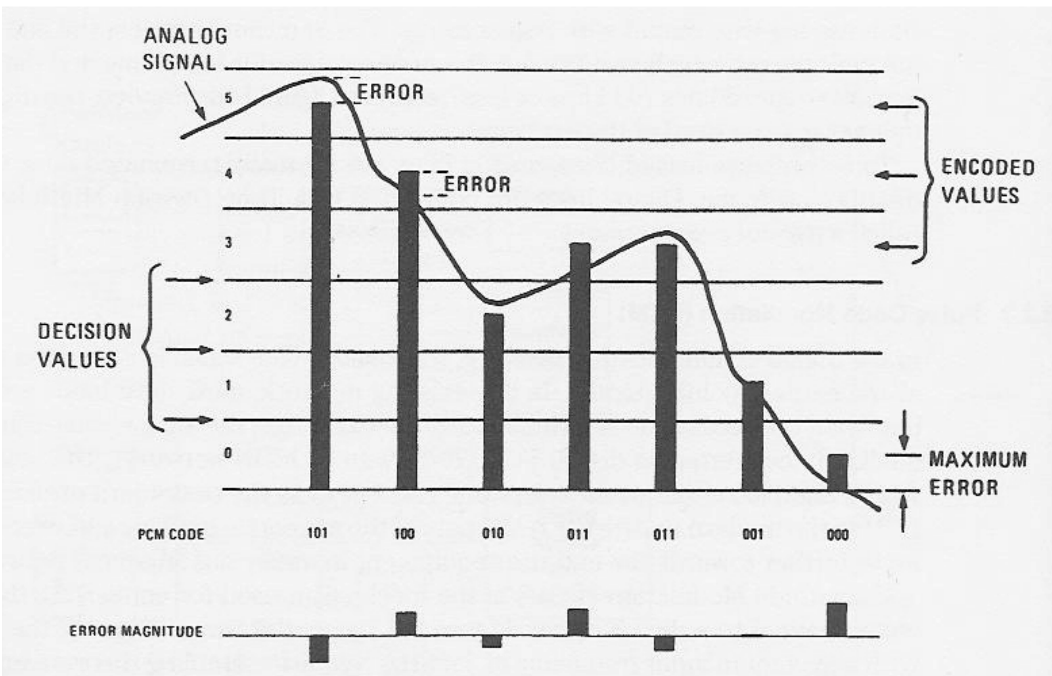

Digitizing a continuously variable analogue signal always results in quantization error or noise, because the discrete digital values never exactly match the analogue signal

A conventional voice telephone signal has a spectrum contained within a 4 kHz bandwidth, more specifically 300 Hz to 3.4 kHz. With PCM encoding, the analogue telephone signal is sampled 8000 times a second and each sample is represented by an 8-bit word. This gives a bit rate of 8000 x 8 or 64 kbit/s.

[2] Counting the errors and measuring the quality of these digital links was the job of South Queensferry's family of BER Testers or Error Performance Analyzers described in Chapter 9.

As mentioned earlier, 8 bits only gives 256 levels or +/- 128 quantized levels. This works well with a high level signal, but low level signals will sound very "grainy" and unnatural. To overcome this, the telephone signal is compressed so that the higher level signals use fewer quantized levels and the low level signal is encoded as if it were passing through a 12-bit rather than 8-bit quantizer as shown here:

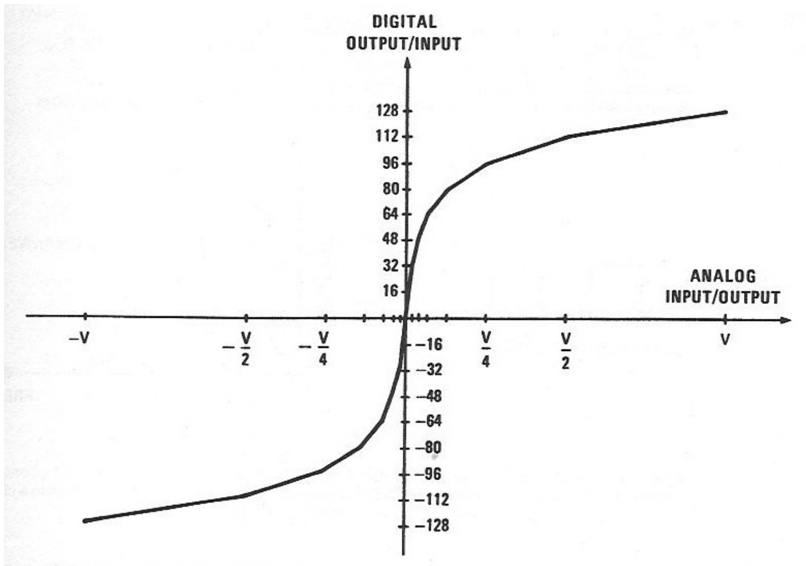

To reconstruct the analogue signal at the receiving end, the reverse process of expansion is necessary. The process of compression and expansion is called "companding", and is one of the main reasons it is necessary to measure the PCM system since the compression and expansion must be an exact inverse or the signal will be distorted at the receiver. Various methods were used for evaluating this non-linearity and measuring the resulting quantization noise or distortion. Some measurements were specified using a sinusoidal tone or multiple tones and some using a random noise stimulus.

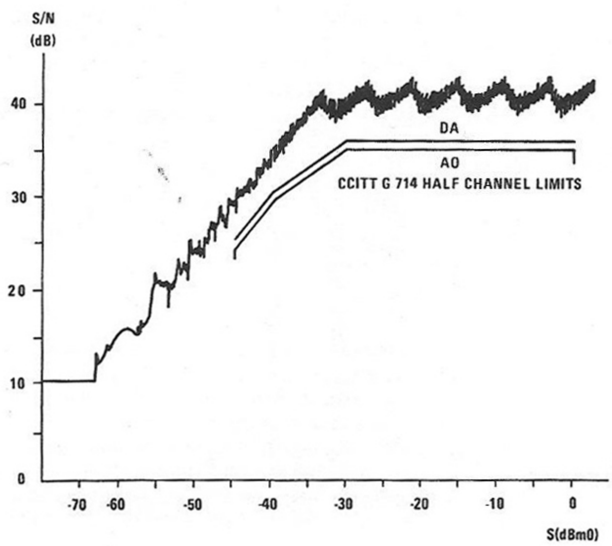

For 8-bit encoding, A-law PCM encoding produces a S/N ratio vs.signal level graph like this. For low level signals, the coder operates as if it were a 12-bit rather then 8-bit encoder. For higher levels, the compander maintains a constant S/N ratio by progressively compressing the signal. A-law is used internationally, North America uses µ-law encoding which is slightly different

A second need for measurements arises because of the sampling process. Because the telephone signal is being sampled at 8 kHz, it is important that no analogue signals are present above 4 kHz (Nyquist's Sampling Theorem). Likewise at the receiver, all signals above 4 kHz must be suppressed. This means there needs to be tight filtering at the coder and at the decoder, restricting the bandwidth to 3.4 kHz. These filters must do the rejection of unwanted signals while not interfering with the telephone signal. It is necessary to measure the bandwidth or frequency response of these filters.

Two main standards evolved during the 1960s, one in North America and the other in Europe. The American standard as already described, combined 24 analogue channels into a PCM bit stream at 1.5 Mb/s (known as T1), whereas the European standard, also used internationally later, combined 30 telephone channels into a PCM bit stream at 2 Mb/s (known as E1). These were the lowest levels of the digital multiplex and are sometimes referred to as the primary multiplex. The two standards also used slightly different "companders" (A-law internationally and μ-law in North America), so different test equipment was required.

In the early days of PCM, it was just point-to-point links, so tests could be done from analogue telephone input to telephone output with the transmitter and receiver connected back-to-back through the digital T1 line. This was the way "PCM channel banks" were measured. By the early 1970s, international standards for the performance of PCM coders and decoders ("codecs" for short) had been developed and the methods for making measurements (CCITT[3] Recommendations G.711, G712, first ratified in 1972).

The next stage of development was the evolution towards a fully digital network, which included digital telephone switches and digital transmission systems. In time, this Integrated Digital Network (IDN) would mean the analogue telephone signal only existed at the edges of the network. PCM encoders and decoders then became part of the digital telephone switch which terminated the telephone subscriber's line. These were called line cards, with maybe 2, 4, or 8 lines terminating on the card.

By the early 1980s, millions of these cards were being manufactured every year as the digital revolution gathered pace. The need was then to measure the line cards Analogue to Digital (A-D) or Digital to Analogue (D-A) to ensure they could inter-work correctly with any other line cards elsewhere in the network.

These measurements were specified in later CCITT Recommendations G.714 and G.715, and are referred to as "half-channel" measurements. These are much more complex, as measurements at the digital line interface require extracting the bit pattern from the time slot corresponding to the telephone channel under test, decoding the PCM signal with an "ideal" decoder and analysing the equivalent analogue parameters. Likewise on the D-A measurement, the test signal needs to be accurately PCM encoded and inserted into the appropriate timeslot in the framed digital line signal.

This was the measurement challenge addressed by South Queensferry's PCM test sets in the 1970s and 80s.

[3] International Consultative Committee on Telephone and Telegraph which in later years became ITU-T (International Telecommunications Union – Telecom Standards)

HP 3779A/B Primary Multiplex Analyzer

Not long after the first CCITT PCM standards were published in 1972, the product development group at South Queensferry started to investigate how to build a PCM test set. At that time, work was nearly completed on the Division's first pattern generator and error detector, the 150 Mb/s 3760A and 3761A, so no doubt the product group would already be aware of the PCM test requirement. Rod May remembers being appointed Project Leader in November 1973, marking the start of development. Rod had joined HP South Queensferry in 1970, working first as a marketing engineer on the Dynamic Signal Analysis products, followed by a period in the Division's System Integration Centre where he gained valuable experience in computer controlled measurement systems.

The CCITT G.712 standard of 1972 relates to performance measurements Analogue to Analogue (A-A), assuming the PCM coder and decoder are connected back-to- back at the digital interface. However, HP found that customers were also interested in the half-channel measurements (A-D and D-A) for the reasons given above.

The then Marketing Manager, Finlay Mackenzie, had recruited Gordon Pasque from STC in South Africa:

"*He worked with me for a year before going back to South Africa and the need for a test instrument like the 3779 was largely based on his input, having worked in STC looking at the R&D and production requirements.*"

There were discussions also with engineers at the British Post Office and in Norway who made a strong case for the digital measurements to ensure full inter-working between different equipment and across network boundaries. The 3779 would incorporate all these measurements.

Rod May and Peter Huckett from marketing did much of the market research for the new product in Europe and North America, and by late 1974, the definition and block diagram were complete. Work started on a marketing prototype to test the concepts. The list of measurements (right) was extensive and required a lot of analogue and digital hardware, as well as processor hardware and software for controlling the measurements.

Measurement capability

Measurements	A-A	A-D	D-A	E-E
Gain	●	●	●	●
High accuracy gain	●			
Gain using peak codes		●		
Digital mW gain			●	
Gain vs frequency	●	●	●	●
Gain vs level using noise (3779A only)	●			●
Gain vs level using tone	●	●	●	●
Gain vs level using peak codes		●		
Gain vs level using sync 2 kHz			●	
Pedestal (coder offset)		●		
Idle channel noise psophometric (3779A only)	●	●	●	●
Idle channel noise C-message (3779B only)	●	●	●	●
Idle channel noise 3 kHz flat	●	●	●	●
Idle channel noise selective	●	●	●	●
Noise with tone	●			●
Quantizing distortion using tone	●		●	●
Quantizing distortion using noise (3779A only)	●			●
Intelligible crosstalk	●	●	●	●
Intermodulation using two tones	●			●
Intermodulation using four tones (3779B only)	●			●
Discrimination against out-of-band inputs	●			●
Spurious out-of-band outputs	●			●
Spurious in-band outputs	●		●	●
Return loss (Tx and Rx)	●			
Impedance balance (Tx and Rx)	●			
Signal balance	●			●
E & M signalling distortion	●			●
Analog level	●			
Digital level		●		
Remote alarms (3779A only)			●	
Multi-frame alignment (3779A only)			●	
Frame alignment (3779A only)			●	
Local alarms (3779A only)			●	

A-A : analog-analog. A-D : analog-digital. D-A : digital-analog. E-E : end-end.

The marketing prototype was assembled in 1975 and consisted of three instrument cases. One contained the analogue and digital measurement hardware, one the display, printer and power supplies and a third case with the processor which was an Intel 8080 with memory.

The analogue transmitter mainly involved sending single or multiple sinewave signals to measure gain, frequency response, distortion, quantization noise and level. These were generated using an ingenious direct digital waveform synthesizer developed at South Queensferry in the early 1970s for the 3770A telephone line analyzer. This produced accurate and stable signals in 10 Hz steps with low harmonic distortion. Based on a memory look-up table for the sinusoidal shape, the 3779 implementation was operated at a higher clock rate so that simultaneous multiple sinewaves could be generated by interleaving the digital addressing of the memory. Because of the limited size and high cost of semiconductor memory in the 1970s, the designers went to some lengths to reduce memory requirements by segmenting the sine waveform so that the same look-up could be used for the four 90 degree segments. (See Chapter 6 on Telephone Line Analyzers for more details on the 3770 digital synthesizer.)

The output of the digital sinewave generator was a sequence of 12 bit-wide words representing the sinusoidal signal. These could be used to drive a digital to analogue converter for an analogue output, or could be used as the source for the direct digital PCM signal to be inserted into the digital timeslot. An ingenious "digital attenuator" added or subtracted shifted copies of the 12-bit sample to set the level. This process used 20-bit arithmetic. The output was then compressed to give the 8-bit PCM word for insertion in the timeslot.

About half the instrument was taken up with a multitude of analogue filters (active filters) for generating test signals and measuring band-limited noise and distortion products. There was also the need to provide the correct termination for the telephone circuit and set and measure analogue signal levels.

A very big challenge with this instrument was the user control. The table above shows the huge range of measurements possible, each "blob" in the table is a measurement or series of measurements. Early on in the project, the designers decided to use a menu-driven display that would allow the user to select single or multiple measurements in a sequence and set parameters – a novel idea at the time. "*You have a problem and you look for logical solutions with the technology available. And then something comes out of the woodwork*", as Rod May recalled. For the majority of tests, the CCITT Recommendation G.712 specified the measurement settings so these could be programmed into the instrument as default settings. The instrument also stored the CCITT performance limits, so most tests could be run as a PASS/FAIL rather than reporting hundreds of measurement results.

In the mid-1970s, South Queensferry's German competitor Wandel & Goltermann, had a PCM test system consisting of four boxes: separate analogue and digital transmitters and receivers[4]. It would have been very time consuming to conduct all the G.712 measurements using this set-up.

[4] W&G PCMG-1 and PCME-1 plus the PDG-1 and PDA-1

The marketing prototype evaluation was done mainly by Rod May and Peter Huckett and lasted three to six months. Then in 1976 work started on building the lab prototype, which would be the first prototype of the final production instrument. It was a "tour de force" of analogue and digital design and state-of-the-art processor control. Like several products developed at South Queensferry in the 1970s, it was ambitious and ventured into new areas, stretching current technology to the limits. In the case of the 3779 this included the digital signal generation and the processor control.

Around this time, HP in the USA had developed a high-performance 16-bit processor chip using silicon-on-sapphire CMOS, as part of their advanced computer developments. It had 10,000 transistors and could run at clock rates up to 8 MHz with an instruction time of around 1 microsecond. It was well suited to the kind of processing required in the 3779, so the designers chose it instead of a commercial processor. The HP processor was fast enough that some of the PCM generation and measurement could be done at the bit level and the measurement results computed directly, thus removing a source of error. It was not without its problems, however. The "super-fast" rise and fall times of the 12V logic created significant interference with the analogue circuits in the lab prototype, although this was eased when HP produced a 5V version of the processor.

The problem was that no standard software development tools existed for this new processor, so the lab engineers had to develop these first, using the lab's HP 2100 mini-computer. The system software was developed on the lab computer using a home-grown assembler and then transferred using paper tape to the large debugging/control board which sat on top of the prototype instrument, and this was used to track what was going on in the processor. Measurement and self-test code was developed using a high-level language called STAB, acquired from Strathclyde University and extensively modified for the environment and application. High-level programming allowed the 3779's array of measurements to be developed, optimised and maintained far more easily, since the source code could be read by various members of the team and it also provided development aids for the measurement routines. It was probably the first time a high-level language had been used in instrument software development. The processor itself was operated at 4 MHz and addressed 32k of 16-bit ROM as well as 2k of non-volatile memory using magnetic core (shown here), plus some semiconductor RAM.

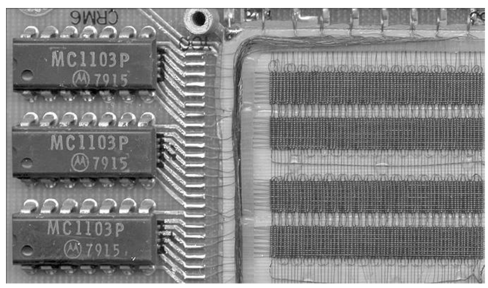

Much of the credit for this highly advanced processor and software development lay with two Division software gurus, Ralph Hodgson and Virgil Marton. Earlier, Ralph had developed the FDM plan algorithms for the 3745 SLMS processor (see Chapter 5) and then worked on the 3779 marketing prototype. After he left HP, he was replaced by Virgil Marton, a Romanian computer scientist who came to the UK in 1968 and joined HP after working for the mainframe computer manufacturer, Burroughs. The instrument was a small mini-computer in itself, and contained more processing power than any other instrument at the time. Following the marketing prototype phase, measurement software development was led by Mark Dykes.

A lot of work was also done on the ergonomics. The mechanical product designer, Dave Leahy, came up with a sloping front panel in cast alloy that housed the CRT display (a small TV tube), the keyboard and connector panel. The CRT hardware and associated electronics had been used by the Loveland Division for its Logic Analyzers. The keyboard had relatively few buttons and relied heavily on navigation through the menu-driven display to set-up measurements. For its time, it was a remarkably innovative design and thought to be the first of its kind. It was an elegant solution to the problem of setting the large number of measurements and parameters in the instrument. In fact it is hard to imagine it would have been possible any other way.

As the instrument contained analogue and digital transmitter-receiver pairs, it was ideally suited to internal loop-back self-testing, and extra signal paths were designed in from the start for this purpose. When the instrument was switched on, a quick test would verify that all was well end-to-end; in case of a fault, diagnostic routines could test hardware progressively to localise faults. This was a boon to development, manufacturing and service of such a complex instrument. Self-test and service aids made up fully one third of the instrument software.

Mark Dykes recalled an amusing incident the first time the analogue receiver was tested:

"*We were testing it using the lab attenuator. Seeing an initial "0.00 dB", Robert Duncan switched in 10 dB and the instrument then showed "10.00 dB" – a plausible result given the hoped-for accuracy. As Robert added more 10 dB steps, bystanders grew increasingly excited as the results showed perfect linearity down to 50 dB! Then Robert became suspicious and dialled-in 51 dB. There was a groan as the reading stayed at "50.00 dB". Still some work to do, then!*"

A full description of the product is beyond the present account, but fortunately two informative articles were published about it in the *HP Journal* for January 1980[5].

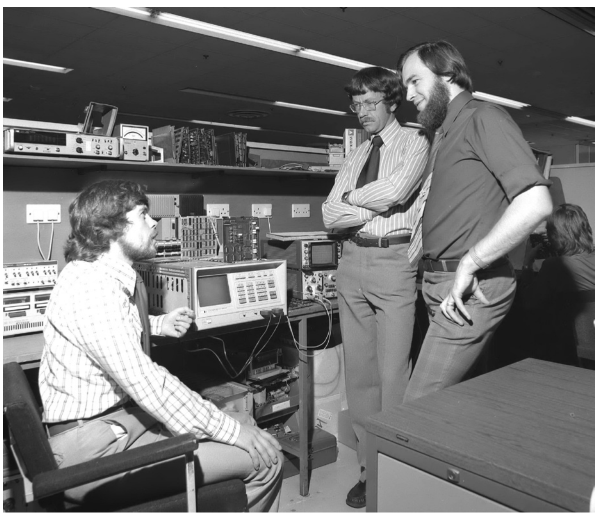

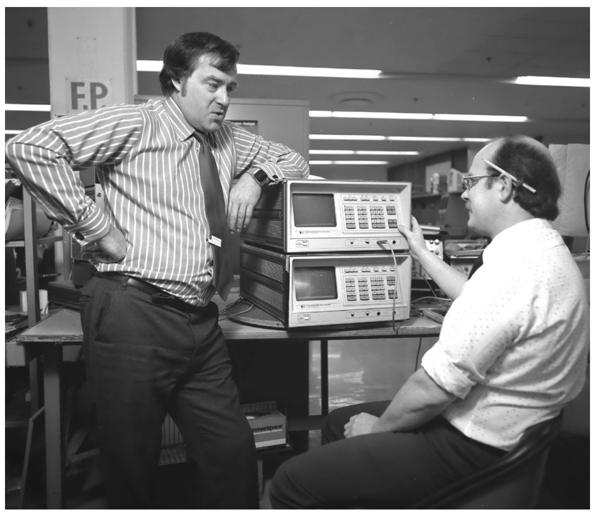

Some of the "PMA" Team – Mark Dykes, Peter Huckett (marketing) and Andy Batham

Peter Carmichael speaks to Alan Gardner checking the 3779 PMA in production

In late 1977, Rod May left HP and Rob Pearson took over as Project Leader through the production prototype phase. With the high software content, bugs emerged even at a late stage. As the first units were being manufactured, EPROMs were being removed,

[5] "*Automated Testing of PCM Communications Equipment with a Single Self-Contained Instrument*" by Pearson, Dykes, Marton, Batham and Bryant; "*Software for an Automatic Primary Multiplex Analyzer*" by Dykes, HP Journal January 1980 pp 3 – 21. http://www.hparchive.com/Journals/HPJ-1980-01.pdf

reprogrammed and replaced on the production line. Alistair Lucas, the Manufacturing Manager, asked Mark Dykes for an assurance that there were no further bugs in the instrument. As Mark recalled, "*I was only able to guarantee the opposite!*"

Two versions of the Primary Multiplex Analyzer were introduced in 1979. The 3779A was designed for the European/International (sometimes referred to as CEPT[6]) market complying with the CCITT Recommendations and operating at the 2 Mb/s (E1) primary rate, while the 3779B operated at 1.5 Mb/s (T1) primary rate for the North American market and conformed to the Bell PCM "companding" standards and performance criteria.

Since the instrument had been designed to test semi-automatically with its measurement sequence and Pass/Fail testing, the project team also designed an access switch, the 3777A Channel Selector which would allow multiple telephone receive and transmit ports to be tested one by one as part of the measurement sequence. This allowed a whole channel bank to be tested automatically without operator intervention. By using multiple switches, up to 256 separate physical telephone ports could be accessed, with the 3779 controlling the access switch through the Hewlett-Packard Interface Bus (HP-IB).

At the other end of the scale, with an eye on the growing codec and line-card test market, the 3779 also had 64 kb/s digital ports for easy interface to single-channel devices.

Finlay Mackenzie recalls taking a production 3779B on an introduction tour in the USA in the late 1970s. Development and production engineers at AT&T's factory in Merrimac Valley (WECO), Boston, were so impressed with the new instrument they wanted to keep it there and then for the production line:

"*The CEO took me into his grand office and told me he was going to phone John Young (HP's CEO) to tell him they had to have it. After some negotiation, I was allowed to continue with the tour round the USA on the understanding the box would go back to Merrimac when I returned home! Afterwards, apparently, he did phone John Young, who commented on a later visit to Queensferry that we sure understood the market requirements.*"

The 3779 certainly had its enthusiasts amongst the world's telecom manufacturers and operators, but it wasn't universal. The 3779 attempted (and largely succeeded) to include just about all the measurements you would ever want to make on a PCM channel bank or line-card, all in one box. With the technology of the day, this resulted in a very large box, one of the biggest single-box instruments ever designed at South Queensferry.

[6] Committee for European Post and Telecommunications

It had over 30 printed-circuit cards which needed a cabinet 9" high and 24" deep, weighing 55lbs (25 kg). It was hardly a portable box. The selling price was substantial too at $24k, while the 3777 access switch cost $4.3k.

In 1981, an upgraded unit was introduced, the 3779C for the CEPT market replacing the 3779A and the 3779D for North America replacing the 3779B. The new instruments had some enhancements including additional half-channel measurements that filled gaps in the measurement capability table. This was possible with the introduction of more powerful digital ICs and better memory chips. The sinewave and noise generators were now both synthesized through ROM look-up and interleaved by means of four bit-slice processors. This enabled for example, noise stimulus and measurements to be made A-D and D-A, the new versions also allowed D-D measurements to be made, satisfying the testing of new digital switches, concentrators and trans-multiplexers which interconnected the analogue and digital networks. However, most of the measurement and processor hardware remained the same. By then, the unit price was around $27k.

The 3779 Primary Multiplex Analyzer, or "The PMA" as it was known in the factory, was an impressive instrument with a certain "erudite" reputation, since the project team included some of the Division's "high-flyers", particularly on the digital/software side. One of these was Mark Dykes, later Project Leader on the 3779C/D. Mark had observed that the PMA's readings were rock-solid on a back-to-back measurement, but varied when testing a PCM system – quite significantly at lower levels. Looking into this, he and Andy Batham discovered the variation was caused by the way the PCM digitizing and quantizing process interacted with the 10 Hz resolution test frequencies available in the PMA. Given the intended accuracy of the PMA and its use in R&D and QA, they thought it a good idea to describe the theoretical cause of this variability.

Mark wrote a technical paper about it entitled "*Overcoming Intrinsic Uncertainties in PCM Channel Measurements*". (Later this was published in a book, see below[7].) In the early 1980s, the Division organised some Symposium tours in Europe and North America, inviting customers to technical presentations and product demonstrations. I remember Mark Dykes, alumnus of Jesus College Cambridge, with a slight frown on his face delivering his "*Intrinsic Uncertainties*" to the bemused customer audience (sounds a bit like Heisenberg again, does it not). The message was "*If you think the PMA is making errors or mistakes, it isn't*". It also may have helped to cement the reputation with the sales organisation, that the Division was the "University of South Queensferry".

Graeme Nelson from product marketing recalled an amusing experience presenting the same paper:

"*I was on a massive six-week tour of North America in 1981, taking over from Mark Dykes for the "Intrinsic Uncertainties" presentation on the second leg of the North American Symposium, ending up in the San Francisco Bay Area. My colleague, Reid Urquhart, had finished his presentation and it was my turn with the PCM paper. Reid watched in amazement as the fairly large conference room got fuller and fuller until there was standing room only and then the doorway got crowded and the corridor beyond. None of us knew exactly why this hapless bunch of telecoms people had turned-up, but they were soon to be recalibrated as we went from 'Quantizing Mechanisms' (OK, still got the*

[7] "*Telecommunications Measurements, Analysis and Instrumentation*" by Dr. Kamilo Feher/Engineers of Hewlett-Packard, Chapter4. Prentice-Hall Inc. 1987, ISBN 0-13-902404-2

audience) to 'Probability Density Functions and Quantizing Gain' (mmmm … more than a few glazed looks) to 'Gain Variation with Start Phase'.

"By this time I KNEW I was losing the audience and tried to up the excitement level by saying 'now here's a very interesting point' and then, in Reid's own words, proceeded to a point of stultifying boredom! By the time I got to 'Gain Error vs. Skew' and 'Bit Stealing' there was an aura of panic in the room. I think the concluding applause was rapturous in its recognition that the end had finally come!

"Later, an AT&T or PacBel 'greybeard' said he was fascinated by the possibility of a +/- 0.004 dB variation being caused by 'Gain Error vs. Skew' and said, 'Hell son, if my guys can lift the phone and yell down it and the guy at the far end can hear it, that's good enough for me!!'"

Sales were reasonable but not spectacular considering the rapidly expanding market in digital switching. Some customers were very keen on the PMA, and Mike Kerr mentioned that Northern Telecom nearly sued HP because of late delivery of this highly desirable product! It was certainly an excellent box for R&D work on channel banks, line cards and codec chips, but would have been too expensive and bulky for significant field use.

In 1980/81, the marketing group estimated the Division had about 25% market share (in dollars) for PCM instruments. Around 8 units/month of each model sold until the mid 1980s and it declined thereafter until the 3779C/D was discontinued in 1994. Between 1979 and 1994, just over 2100 units in total were sold of all four variants, bringing in over $50M. Around 1650 3777A Channel Selectors were sold, adding another $7M.

As mentioned earlier, a big market existed in line cards by the early 1980s, as network digitization gained momentum. Several manufacturers were producing large scale digital telephone switches (BT System X, AT&T 5-ESS, Nortel DMS-100 etc.) handling 50,000 lines or more. Millions of line cards for these switches were in production, and some were being tested with PMAs. Mike Kerr from marketing recalled a visit to the Western Electric factory in Oklahoma which confirmed the scale of production.

South Queensferry toyed with the idea of building a product more specifically for this market, and had discussions with the board tester development team in Loveland Division, Colorado. Although the 3779 PMA with its automatic sequencing was much faster than the manual test sets, it still made the measurements serially. Because of the volume, line card testing needed to take seconds not minutes. The market at that time was dominated by companies such as Teradyne and LTX who built what were dedicated board-test systems with some added functional tests. The measurements were done in parallel and fast – "intrinsic uncertainties" not withstanding!

South Queensferry decided to stick to its traditional market and see what could be done to reduce the size and price of a PCM test set.

3776A/B PCM Terminal Test Set (Unit Preserved NMS T.2010.68)

Nearly half the 30 printed circuit boards in the 3779 were partly or wholly occupied by analogue active filters, mainly for measurement of noise or tones in different bandwidths. This was a lot of hardware so some way of reducing it was the key to a smaller box.

Enter David Dack, one of the more inventive design engineers at South Queensferry in the 1970s. Starting work on the Dynamic Signal Analysis Correlator product family, in 1972 he joined the 3745 SLMS project leading the first microprocessor development in HP in that pioneering instrument. He wanted to build a fax machine before there were fax machines, and he invented the first of South Queensferry's highly successful family of HP-IB extender products. Steeped in software design and digital signal processing, David had been agitating for years that the Division should develop a digital filter IC that could be programmed for many different applications and products. The PCM test set seemed an ideal opportunity, as the filters were all at audio frequencies where the digital filter idea was most applicable at that time[8].

Around 1979, he and another engineer, John McElroy, designed a concept circuit for a programmable digital filter chip. David Dack remembered giving a demo of the prototype to Al Bagley[9], corporate engineering manager:

"*All he could see was this black box connected to a signal analyzer, displaying a Butterworth low-pass filter. I could imagine his brain working, 'What has this idiot dragged me over to see this for?' Then, at the flick of a switch, it was a Butterworth high-pass, then a very sharp notch, then a Psophometric filter, then a C-message, then a 12th order Bessel band-pass, all with very low noise floor and mathematically perfect response. He almost fell off his chair!*

"To his credit, he organised an HP-wide conference, to which I was summoned a few weeks later. After my presentation, I was approached by some slightly annoyed R&D section managers, saying, 'How come you are allowed to have so much fun designing stuff like that? We are too snowed-under with management to have time for that sort of thing.'"

The University of South Queensferry certainly benefited by its distance from Palo Alto!

As John McElroy recalled,

"*The unusual team David Dack pulled together for the project consisted of myself and two university staff, John Mason from Swansea and Alwyn Owens from Bangor. We worked on both digital signal generation and digital signal processing. It was great fun and it led to a two-year secondment to the Loveland Technology Centre (LTC) in Colorado where there was a microchip facility.*"

---

[8] "*Digital Signal Processing in Telephone Channel Measurements and Instrumentation*" by David Dack and Bob Coackley, Chapter 3, *Telecommunications Measurements, Analysis and Instrumentation* by Dr. Kamilo Feher/Engineers of Hewlett-Packard. Prentice-Hall Inc. 1987, ISBN 0-13-902404-2

[9] Al Bagley was noted for the joint development with Len Cutler of HP's famous atomic clock time standards in the 1960s and 70s at the Santa Clara Division in California.

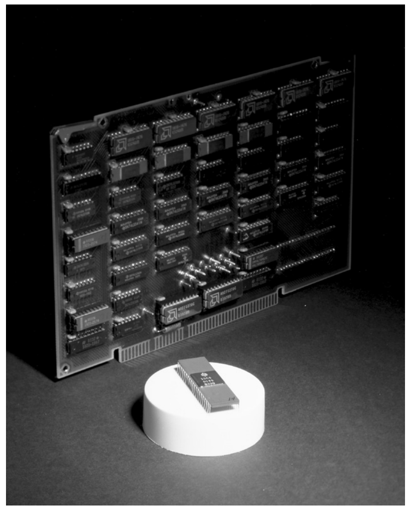

With the help of staff there, John took a course on integrated circuit design and engineered the chip which had eight layers of artwork. The device used a 4 micron silicon NMOS process and had about 10,000 transistors, similar in size to the Intel 8080 microprocessor of the time. It was the first IC that LTC fabricated with a map of Scotland on it! After 18 months, the device was complete and production samples were ready in early 1982.

This was then the key building block for the new 3776 PCM Terminal Test Set. It eliminated a significant amount of circuitry in the old 3779. The new instrument borrowed the digital synthesizer circuits using the bit-slice processors from the 3779, and although it was an analogue and digital measuring instrument, much of the internal electronics was now digital and more compact. The new instrument used a case less than half the volume of the 3779 and weighed around 33lbs (15kg) compared to the 55lbs of its predecessor.

The digital electronics further reduced manufacturing cost as there was very little adjustment required of analogue circuits. The PMA had used a significant amount of auto-calibration, and the 3776 took this even further allowing the use of less tightly specified analogue components and ensuring better long term stability.

The project team of nine engineers and software designers was led initially by Rob Pearson who had previously worked on the 3779 PMA. Later, Andy Batham took over as project leader. Several of the team members had worked on the PMA so there was a lot of latent expertise and knowledge of PCM measurements in the group.

It was limited to an upper analogue frequency of around 4 kHz, whereas the 3779 PMA operated up to 40 kHz on its analogue interfaces. For most uses, this limitation was not a problem and to compensate, the digital filter chip allowed space for the designers to add an optional set of voice-data measurements (group delay, phase jitter, impulse noise etc.) all available through both analogue and digital interfaces – a first.

Another reason for the smaller size was the loss of the CRT display with the menu-driven user interface and measurement sequencing. The 3776 selected from the long list of measurements using a block of keys and LEDs. It was expected that measurements would be done manually one at a time using the preset measurement parameters defined in the standards. One could set level and frequency manually, but it was less convenient than it had been with the numeric keyboard in the PMA. Described by some as "*a bit of a kluge*", it did demonstrate the superiority of the PMA's menu-driven user interface.

Graeme Nelson, from product marketing, recalled an amusing incident when demonstrating the product in California, which highlighted the problems of a complex front-panel:

"*I was showing the box to some representatives from the telephone companies about the time of divestiture in 1984. There was also an old guy from AT&T keeping an eye on things. He was chewing a green cigar and had a face like a tobacco leaf. I was demonstrating the intricacies of making data impairment measurements on a PCM voice channel, when the old guy butted in. 'Is that box meant for the craft? 'Cos what you need to understand is that the techs have an IQ of about 13, that's somewhere between a rock and a head of lettuce. They need three buttons, Start, Stop, and Help I'm f*****g stupid.' This was followed by silence. My pitch went as flat as a deflated Whoopee Cushion!*"

It was possible to download measurement sequences and store them in non-volatile memory, so the instrument could be run automatically as a stand-alone unit. Some customers didn't like all the control built into the PMA and preferred the 3776 which they could program using an external controller. Thus the two instruments served slightly different needs, and both were maintained in the catalogue until 1994.

Like the original PMA, the 3776 was sold in two versions. The 3776A applied to the CEPT market with the 2 Mb/s (E1) primary multiplex, while the B version handled the North American standards and the T1 digital interface. Introduced in early 1984, the 3776 sold for around $11.5k, half the price of the PMA. Both the A and B version sold around 1450 units over the 10 year production life, but with markedly different sales profiles. The European 3776A sold around 12/month in the early years declining to 5/month in the 1990s. The North American 3776B sold over 1100 units between 1984 and 1986, presumably as part of a large deal to supply AT&T with instruments for a remote test system called DARTS[10]. Thereafter sales dropped to around 4/month and virtually none in the 1990s. Overall the combined sales were 2900 units with revenue of around $32M.

All the PCM test products and the 3777A switch were withdrawn in 1994 and not replaced. Interestingly, the Modular Hornet platform used for new products in the 1990s, with its menu-driven flat screen display and soft keys, would have been ideal for a new PCM tester. However, market demand just wasn't there any more. By then the technology had moved on to more sophisticated voice coding schemes with greater bandwidth efficiency than PCM. Abandoning a product area was always a bit difficult as Finlay Mackenzie recalled:

"*It was difficult to leave PCM measurements behind because of pressure from the field sales force. I remember Bill Terry (Test and Measurement VP) coming to see me after a visit to Canada. He said, 'The field are wanting you to do another 3779,' and bending my ear on it. However, I felt we had to move on, and he accepted this.*"

The HP 3776B preserved in the National Collection was built in mid-1984 and was the 426th unit in the production run.

Acknowledgements

Former HP employees Mike Hurst, John McElroy, David Dack, Rod May, Mark Dykes, Graeme Nelson, Mike Kerr and Finlay Mackenzie helped with this chapter.

---

[10] Digital Access Remote Test System similar to AT&T's earlier Switched Access Remote Test System (SARTS), see Chapter 7.

9

Chapter Nine

Digital Transmission Analyzers 1970 – 1990

The next two chapters tell the story of the most important family of products designed and manufactured at South Queensferry, spanning a period of 35 years. Starting with the earliest developments in digital telecommunications in the late 1960s, through to the ultimate development of high-capacity fibre-optic transmission systems, the family of Queensferry's Digital Transmission Analyzers, or Bit Error Rate Test Sets, mirrored this evolution towards the modern global digital communication network. The HP analyzers made a significant contribution to that development.

Because of the long period and the large number of products, I have divided this topic into two chapters. The first covers the early evolution of the product line and the classic products developed up to around 1990. This turned out to be a pivotal year for the Division. A new instrument platform was developed codenamed "Hornet" which used Surface Mount Technology (SMT) and gate-array integrated circuits for the electronics. From this time onwards, the digital transmission market was dominated by fibre optics and the new synchronous digital hierarchies (SONET and SDH). South Queensferry exploited this "new wave" to great effect as described in Chapter 10 that follows.

The technical background to this important area is covered in several chapters in "*Communications Network Test & Measurement Handbook*" by Coombs and Coombs (McGraw Hill, 1998)[1], and this is referred to in the pages that follow. Several engineers from HP South Queensferry contributed to this Handbook.

---

[1] "*Communications Network Test & Measurement Handbook*" by Coombs and Coombs, McGraw-Hill (1998) ISBN 0-07-012617-8.

Measurement Background

The fundamental measure of performance or quality in digital systems is the probability of any stored or transmitted bit being received in error. With the latest equipment, the probabilities are very low, of the order of 10^{-12} or less. In order to simulate live data, a system is tested using a Pseudo Random Binary Sequence (PRBS) pattern which appears to the system under test as a random data stream of 1s and 0s, but is in fact deterministic so the measurement receiver knows exactly what should be received and therefore can detect any errors in the data.

To test a digital transmission system, a pattern generator is connected to the input of the system under test, and an error detector is connected at the output

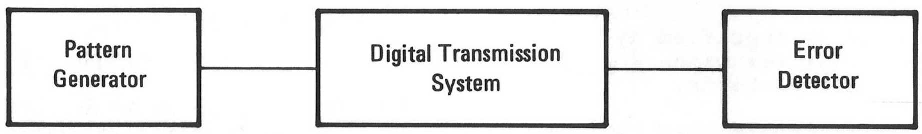

The pattern generator generates a PRBS using a shift register with specific feedback connections to an Exclusive-OR gate to create a maximum sequence length. The number of shift registers determines the length of the PRBS and this is defined in international telecom standards (originally CCITT and now ITU-T) for each of the standard transmission bit rates as shown below.

Bit rate (f_b), kbps	Sequence Length (n)	Polynomial	Spectral Line (f_b/n), Hz
1,544	2^{15}–1 bits	$D_{15} + D_{14} + 1 = 0$	47.1
2,048	2^{15}–1 bits	$D_{15} + D_{14} + 1 = 0$	62.5
34,368	2^{23}–1 bits	$D_{23} + D_{18} + 1 = 0$	4.1
44,736	2^{15}–1 bits	$D_{15} + D_{14} + 1 = 0$	1365.3
139,264	2^{23}–1 bits	$D_{23} + D_{18} + 1 = 0$	16.6

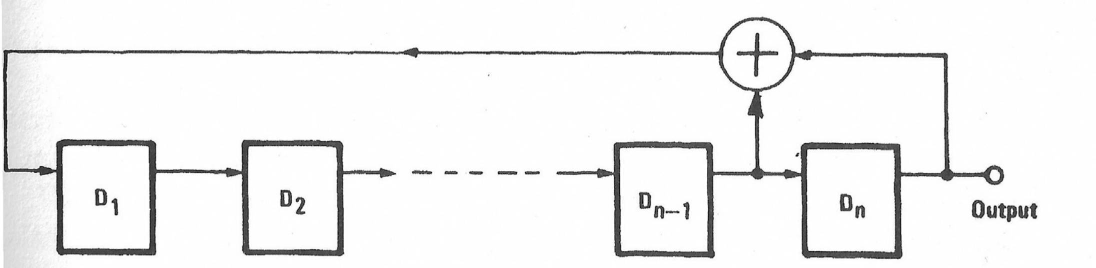

The error detector has an exactly similar PRBS generator which is synchronized[2] with the incoming pattern and then an Exclusive-OR gate compares the incoming stream with the reference PRBS to detect errors:

[2] To obtain pattern synchronization in a closed loop error detector, the feedback loop on the reference pattern generator is opened temporarily and a sample of the input bit pattern is fed into the shift register. The loop is

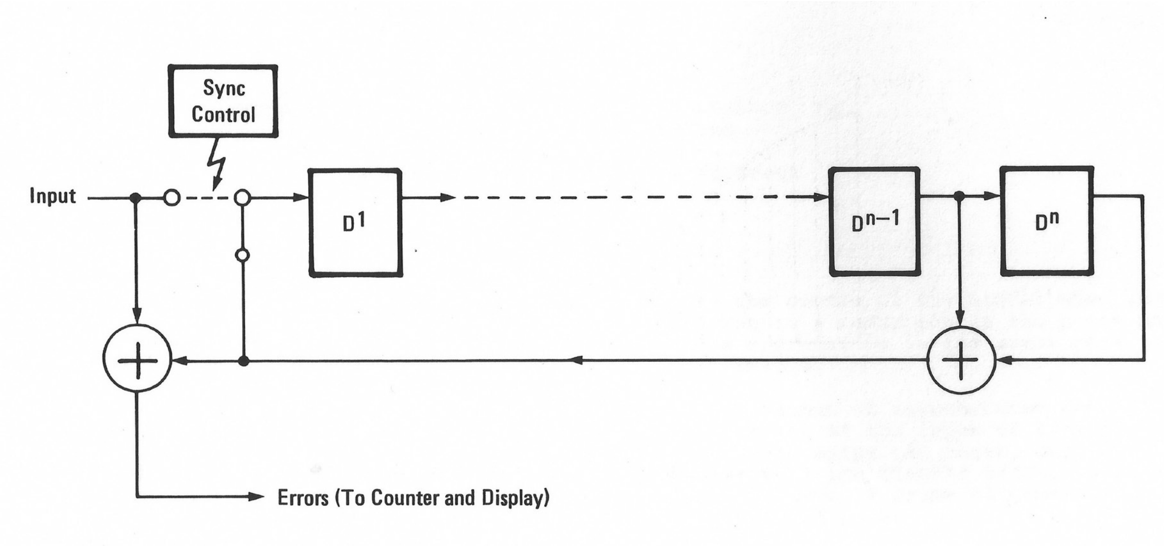

The error count is then accumulated over a period of time and the Bit Error Ratio (BER) is calculated by dividing the error count by the total number of bits received. On the earlier instruments designed in the 1970s, the total bits were simply accumulated with a decade divider to give totals of 10^6, 10^8 and so on and this defined the measurement period.

From the 1980s onwards with more sophisticated processor-based electronics, the measurement time was taken directly in seconds, minutes or hours as this conformed to the various quality of service specifications which require statistical analysis over a given time. Reflecting this more complex analysis of error statistics, the test sets became known as Digital Transmission Analyzers rather than Pattern Generators and Error Detectors or Bit Error Rate Testers (BERTs). This analysis is described further in Chapter 26 of the Handbook mentioned earlier.

In addition to generating and receiving PRBS, all the HP analyzers could also send and receive a user-defined repetitive word pattern, usually of 8 or 16 bits. This is useful for checking the response of systems to specific patterns rather than PRBS.

Fundamental to digital communications systems using Time Division Multiplexing (TDM) is the need to incorporate framing bits in the data stream. These framing bits (sometimes also referred to as overhead bits) are inserted at the sending end at specific points in the data stream and allow the far-end receiving station to lock on correctly to the start of a sequence of random traffic bits. This allows the receiver to extract (or de-multiplex) specific channels or groups of channels embedded in the high-speed stream. The framing bits can also carry alarm and management information. This is described in more detail in Chapter 3 of the Handbook. By the late 1980s, the HP family of Digital Transmission Analyzers (using new programmable gate-array chips) began to incorporate framing as well as simple PRBS so they could check a greater range of specifications in digital transmission systems.

Up until around 1990, all the HP BER testers were designed to operate with the first generation digital transmission systems known as the Plesiochronous Digital Hierarchy or PDH systems. These had multiplex rates of 2/8/34/140 Mb/s. Again Chapter 3 in the

---

then closed and if the output of the error detector is 0.5 this indicates loss of synchronization and the process is repeated until synchronization is obtained.

Handbook has more details. In the 1990s, new transmission equipment increasingly used Synchronous Multiplexing, usually operating over optical fibre, known as the Synchronous Digital Hierarchy (SDH) or Synchronous Optical Network (SONET) in North America. The background to this is discussed in the next chapter. Products developed in the first 20 years were all designed for PDH systems.

Finally, any digital receiver needs to have a clock signal synchronized with the incoming data stream so that it can sample the data at exactly the right instant to determine whether a "1" or "0" is present. Pattern generators and error detectors can operate in two modes. In the simple mode, the pattern generator sends a binary data stream on one output and simultaneously a synchronized clock signal on another output. Likewise, the error detector receives the binary data through one input and the clock through another input port. This is known as "Binary Data and Clock" and is often used for local measurements on the bench.

With installed digital transmission systems however, the receiver may be many miles from the transmitter and is usually connected with a single coaxial cable, radio link or optical fibre. In this case, the receiver must recreate a synchronizing clock from the incoming live data stream. This is usually done with an injection-locked oscillator or phase-locked loop tuned close to the nominal system bit rate. The problem is that if there were not enough transitions on the incoming data, there would not be enough content for the clock recovery to lock onto and it would drift off frequency. In an extreme case, continuous "1s" or "0s" would provide no transitions at all.

In telecom systems, this is solved by using a standardised interface code whereby the binary data is encoded to ensure it has a minimum number of transitions so the clock recovery works correctly under any data pattern. The receiver then needs to remove the coding before the data can be sampled. Nearly all HP's digital transmission analyzers have this interface coding/decoding function built in accordance with international hierarchical interface standards (CCITT, later ITU-T, G.703)[3].

A Product Line is Born – The 3760A/3761A Pattern Generator and Error Detector

In the mid-1960s, the new R&D lab at South Queensferry developed probably the world's first commercial digital noise generator, the 3722A, which was introduced in 1967 (see Chapter 13). It used PRBS sequences to generate pseudo-random noise for structural analysis and other applications. It meant that expertise already existed on the design of digital sequence generators using shift registers.

In 1969, work started on an ambitious project to design the first high-speed pattern generator and error detector. The year before, Motorola had introduced a family of high-speed Emitter-Coupled Logic (ECL) in IC form referred to as MECL3. The Motorola MC1600 series was their third generation of ECL and was the ultimate in high-speed monolithic technology, with switching times of 1 nanosecond. The high speed meant high

[3] CCITT stands for International Consultative Committee for Telephone and Telegraph which later became ITU-T standing for International Telecommunications Union – Telecom Standards. G.703 "Physical/Electrical Characteristics for Hierarchical Digital Interfaces", which defines parameters such as clock rate, pulse shape and coding.

power consumption and the devices required careful circuit design, taking care of propagation delay, to achieve the best performance, so it was a challenge for the design team using this new technology.

The new product would be the 3760A pattern generator and 3761A error detector. Initially the top speed was to be 120 Mb/s but during development this was raised to 150 Mb/s following discussions with potential customers such as British Post Office Research Labs at Dollis Hill and Bell Labs in New Jersey, USA. The BPO (now BT), and companies like GEC, STC and Plessey, were developing high-capacity coaxial cable transmission systems at 120 Mb/s and 140 Mb/s, while in North America, interest centred round the LD-4 or TM-4 transmission systems operating at 274 Mb/s (2 x 137 Mb/s). Around 1970, these bit rates were very high in the field of telecommunications. The only significant digital telecom applications were 1.5 and 2 Mb/s trunks between exchanges (see Chapter 8 on PCM Testers). The new project was led by Lance Mills, a young engineer who joined HP at Bedford, initially as a vacation student working on the control loops in the 5090B Droitwich Frequency Reference Receiver. He and Finlay Mackenzie travelled extensively in Europe to define the new instrument.

Finlay, one of the pioneers of the Microwave Link Analyzer, had transferred to product marketing to head up the new investigation:

"*We started defining the 3760 after a prompt from Post Office research Labs at Dollis Hill to give them some digital pattern generation capability. We also did market research with GEC in the UK, the French research operation CNET, and Italian manufacturer Telettra. These customers also wanted error detection and we became aware that the PRBS (being deterministic) could be used as a remote reference for bit-by-bit error checking at the receiver. I remember saying to Lance Mills that because of the noise generator work, we were fortunate that the expertise and insight already existed at South Queensferry.*"

Finlay also recalled that about this time the Division got its first lesson in the need to be focussed. Dave Packard said South Queensferry should concentrate on telecom test following the success of the MLA, and give up the Dynamic Signal Analysis (DSA) work on Noise Generators and Correlation. Ironically, without that knowledge the Division would not have picked up on the BERT business as quickly.

Raising the bit rate to 150 Mb/s increased the design challenge as engineers Tom Crawford and Ivan Young recalled. Tom had joined HP in 1966 and worked initially on the MLA. In 1967, he left HP to do a PhD at Heriot Watt University in applications for digital noise analysis using PRBS generators. Although not directly related, this research led to a better understanding of the properties of the PRBS, and that it could be generated by interleaving lower-rate PRBS streams.

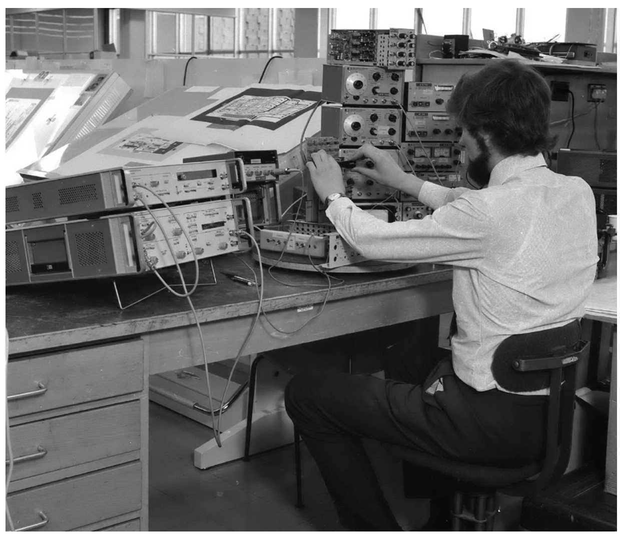

James Robertson works on high-speed circuits

Quite a lot of work had been done on the 3760A pattern generator in which the MECL3 logic ran at full rate up to 150 Mb/s.

This turned out to be problematic. Although MECL3 was very fast and could probably be clocked at 200 to 300 MHz in a simple binary divider, the design was much more challenging in a synchronous logic circuit such as a feedback shift register. In the late 1960s, complex logic design was relatively new as engineers began to understand the problems caused by timing difference between the data and the clock. At 150 Mb/s this was at the edge of what could be done. One can get an idea of the problems in that a 20 cm length of track on a printed circuit board might have a propagation delay of 1 nanosecond, which was about the same as the switching time of a MECL3 gate. The delays in the logic and the interconnections all needed to be considered in the design. In the 3760A, Ivan recalled that the generator would appear to work, but as the clock rate was increased towards 150 Mb/s the PRBS would sometimes lose lock. The fact that the sequence generator logic was spread across two separate printed-circuit boards, may have made the propagation delay problems more severe.

When Tom returned to HP in 1971 he worked on the design of the 3761A error detector, and in this case the reference PRBS generator and synchronising logic all worked at half-rate with interleaving, greatly reducing the design difficulties.

Eventually the problems were resolved and the system was introduced in 1973 with a price of $11.5k. Around 300 units were sold over a seven year period with total revenue of $3.5M. Because it predated digital telecom standards it had only binary data and clock interfaces, and interestingly provided a wide range of PRBS patterns (rather like the 3722A Noise Generator mentioned earlier, that may have been its inspiration). In a way, it was a general purpose parametric test set rather than a telecom specific tester.

The generator had quite an elaborate arrangement for altering the timing between clock and data using fixed and variable delay lines and there was a wide range of output voltage levels and D.C. offsets on clock and data. The pattern generator had a particularly good output waveform with a rise time of less than 1.4 nanoseconds and small overshoot. This was achieved using some proprietary high-speed differential pair transistors from HP Labs in the USA. An article about the design of the 3760A and 3761A, written by the South Queensferry engineers, was published in the November 1973 issue of the *HP Journal*[4].

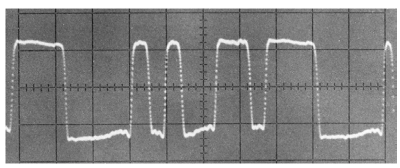

For several years it was one of the fastest BER testers in the world and was a valuable tool for engineers designing the new generation of digital communication equipment. The design team gained very valuable experience which fed through onto the next product developments.

[4] "A High-Speed Pattern Generator and Error Detector for Testing Digital Systems" by Crawford, Robertson, Stinson and Young, HP Journal November 1973, pp 16 – 24. http://www.hparchive.com/Journals/HPJ-1973-11.pdf

The 3780A Pattern Generator and Error Detector (Unit Preserved NMS T.2010.66)

The next product in the BERT family was the 3780A, introduced in 1975. This had a top rate of 50 Mb/s and was therefore less of a technical challenge. It used Motorola's recently introduced MECL10k, which had a switching time of 2 nanoseconds and much lower power consumption.

To define the product, Finlay Mackenzie and Ivan Young visited GTE Marelli in Italy, a company that had previously developed a competitive Microwave Link Analyzer. They had built a BERT for their own use, but after the MLA experience did not want to be in the test equipment business. The Italians got in touch after they became aware of the 3760/61 BER tester. Finlay recalled that they provided South Queensferry with an excellent outline specification for the 3780A, guaranteeing its success. Design work started in 1973 and a lab prototype was completed in 1974.

The 3780A is interesting for a number of reasons. It was the first unit to combine the pattern generator and error detector in a single portable package, similar to HP's 1700 series portable oscilloscopes at the time. It was the first unit to incorporate clock recovery from standard telecom coded data streams (according to CCITT G.703), so in effect it was the first dedicated field portable transmission test set for the new digital communications system, operating up to 34 Mb/s in the international market and 45 Mb/s in North America. Ivan Young recalled that there was some debate as to how certain code errors should be interpreted as binary errors after decoding in the receiver (known as error extension). This was before some details had been tied down by international standards. Over the years a number of different option versions were produced to suit various markets and applications, which extended beyond telecom into military and general purpose measurements.

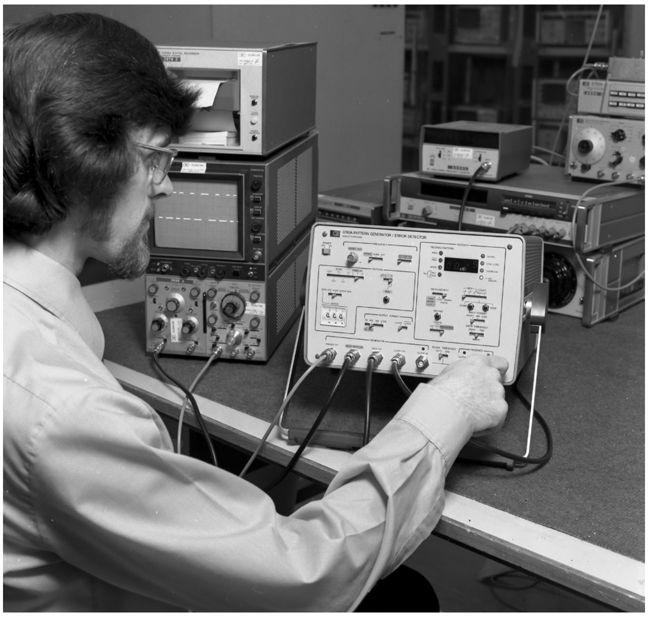

Introduced in 1975, it was one of the factory's most successful products, with a production life of nearly sixteen years during which time over 4500 units were shipped giving revenue of $38M. Its typical selling price was around $8.5k. The design of the 3780A is described in an interesting article in the March 1976 issue of the *HP Journal*[5]. It describes the interface coding/decoding and how clock recovery was done. In this photo, Rob Pearson checks the operation of the new 3780A

Although some development work was started in the early 1980s to replace it (3769A and later 37800A with an interface plug-in), in the end the 3780A was never replaced. This was possibly a missed opportunity, as some of the market served by the 3780A (particularly

[5] "A 50 Mb/s Pattern generator and Error Detector for Evaluating Digital Communications System Performance" by Young, Pearson and Scott, HP Journal March 1976 pp 18 – 24. http://www.hparchive.com/Journals/HPJ-1976-03.pdf

military applications with RS-232 interfaces) was taken by a North American competitor, Telecom Techniques Corporation (TTC), with its FireBERD product. Part of the FireBERD's success was the semi-modular design allowing customers flexibility in specification.

In response, South Queensferry started development in the late 1980s of a highly flexible pattern generation and detection product codenamed ALBERT, based on the industry-standard VXI card cage, using the latest programmable logic, signal processing and software developments with a top rate of 20 Mb/s. (See Appendix 6 for more details.) While it could do the job technically, the customers of the 3780A and the FireBERD wanted a field portable unit, not a card cage. ALBERT was cancelled. The Division inadvertently gave TTC a leg-up in the market and helped to create a formidable competitor in the 1990s. However, on a positive note, this experience was the catalyst for developing the highly successful Modular Hornet family discussed in the next chapter, so every cloud has a silver lining.

3762A Pattern Generator and 3763A Error Detector (Unit in Preservation)

In late 1972, as design work on the 3760A/61A came to an end, the R&D team started an investigation into a future 300 to 600 Mb/s test set, using thin-film circuit technology. This high-speed proposal would use MECL3, but with sub-rate multiplexing as used in the 3761A error detector described above. By June 1974, the upper limit had been set at 300 Mb/s, and the team had built a breadboard prototype. Market investigations then suggested in the short term it would be better to concentrate on a lower cost system with a top rate of 150 Mb/s, but still using the multiplex method of generation so a lower-cost logic family, MECL10k, could be used. As mentioned earlier, this also relaxed design constraints on the logic. One of the project engineers, Bob Thomson, had done a PhD at Heriot-Watt University on the properties and generation of PRBS sequences which was perfect timing for this investigation. Work started on a lab prototype, to be completed in 1975.

This was the next member of the BERT family. With a top bit rate of 150 Mb/s, the 3762A/63A was intended to replace the 3760A/61A, and with a number of feature and performance enhancements. By then, the first set of digital telecom standards had evolved, and the new product took advantage of that with coded interfaces, clock recovery and industry standard PRBS patterns and also automatic cable equalisation at the receiver input. It was equally at home testing new high-speed transmission systems in the field as it was on the development bench. It was an ideal R&D tool with great flexibility at the electrical interfaces as well as optional burst-mode capability for satellite testing.

Although its user interface was over-complicated for field use, it was nevertheless regularly used for installation and maintenance of the new 120 and 140 Mb/s line systems, and the British Post Office (BT) allocated it two pairs of PO Tester numbers, meaning that it was the standard instrument for use on their network in the early 1980s for both coaxial cable and new digital radio systems. For a number of years it was the most versatile test set for 140 Mb/s transmission.

Technically the product was interesting as it was the first HP pattern generator to use half-rate multiplexing whereby the sequence generators ran at half the clock rate and were multiplexed up to the full rate at the output. The product had a number of innovations. The receiver used a novel clock recovery circuit invented by Tom Crawford, an injection-locked oscillator formed of NAND gates and a tuned circuit or delay line. It had a wide bandwidth and was later used in South Queensferry's pioneering work on jitter measurements. At 140 Mb/s the new interface code was called CMI (Coded Mark Inversion). The 3762/63 was one of the first instruments to incorporate this. One of the project engineers, Peter Scott, was granted a patent for a CMI coder/decoder using delay lines which had the advantage of low intrinsic jitter[6].

It was also the first BER tester to use the new thin-film hybrid technology developed at the factory in the mid-1970s. (The other products using it at that time were the 3790 140 MHz Microwave Link Analyzer and the 3747 90 MHz Selective Level Measuring Set). The hybrid circuits were used for the 150 Mb/s outputs and gave a very clean waveform with around 1 nano-second rise and fall times, as shown in this photo of the output of the preserved 3762A operating at 140 Mb/s, measured using an HP Infinium sampling oscilloscope.

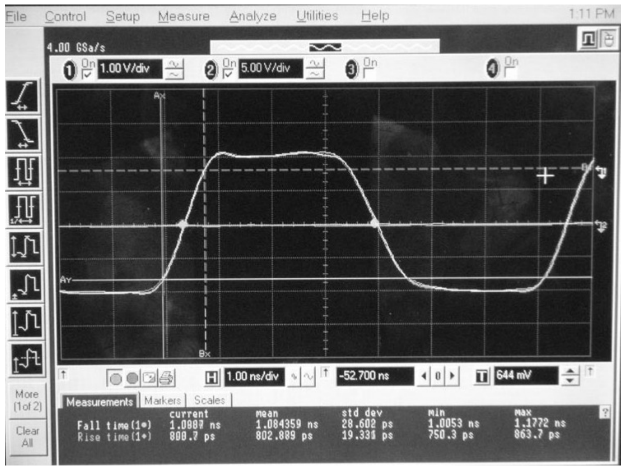

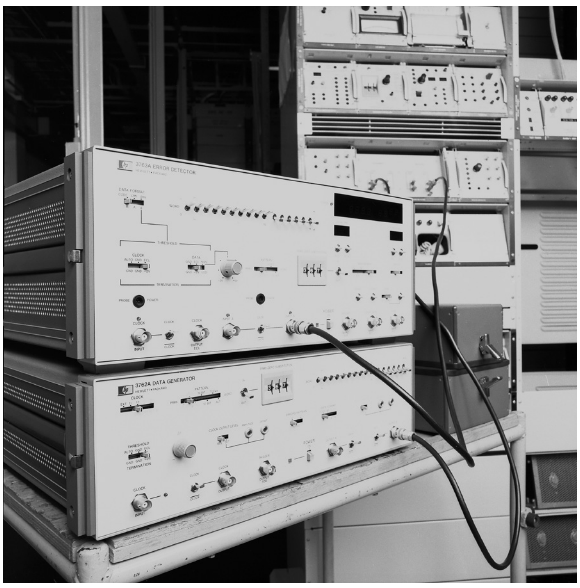

Introduced in 1977, the 3762A/63A system was in production for about 10 years and over 1100 pairs were shipped at an average system price of around $17k, total revenue $20M.

The British Post Office (later BT) bought a good number of these units and it had a type approval number, P.O. Tester 273/274, which made it easier to order within the organisation. This photo of the instruments was taken in a telephone exchange in Portsmouth in 1977 (note the transport covers alongside).

3781A/B Pattern Generator and 3782A/B Error Detector

The three systems from the 1970s, just described, represent the first generation of pattern generators and error detectors. Another system introduced in 1980 also belongs to this group. It was the 3781A/B Pattern Generator and the 3782A/B Error Detector, the A version (shown here) being for the international market and the B version for North America. With a top rate of 50 Mb/s it was in effect the earlier 3780A split into a separate

[6] US Patent 4,189,621 "CMI Encoder", Inventor: Peter M Scott, granted February 19th 1980.

transmitter and receiver. It was more adapted to the specific telecom application and was the first HP tester with measurement gating periods of seconds, minutes and hours rather than total bit count, although it could do that too. Another new feature was an input on the pattern generator that allowed the clock source to be modulated so calibrated amounts of timing jitter could be introduced to check network equipment tolerance. This was Queensferry's first jitter test, covered much more comprehensively by the 3785A/B described next. Priced at around $15k per pair it was a fairly expensive solution. Neither version was a big seller despite the rapidly growing market, managing about 930 units of the A model and 1320 of the B model over 10 years. Total revenue was a respectable $30M.

The final product from this first generation error analyzers was the **3783A Frame Analyzer** for 2 Mb/s. When introduced in March 1980, mature sales were forecast at what turned out to be an unrealistic 50 units/month. A small box, its measurement capability was too limited for the price of $1700 HP wanted to charge, and it was a complete flop selling only 55 units in total!

3785A/B Jitter Test Set (Unit Preserved NMS T.2010.67)

As discussed earlier, all digital systems require a reference synchronized clock in the receiver to sample the data stream at the correct instant to determine if a "1" or "0" is present. If either the data stream or the reference clock (or both) is moving around in time due to instantaneous variations in phase or frequency then there could be problems trying to sample the data correctly and bit errors will result. This variation is called Jitter.

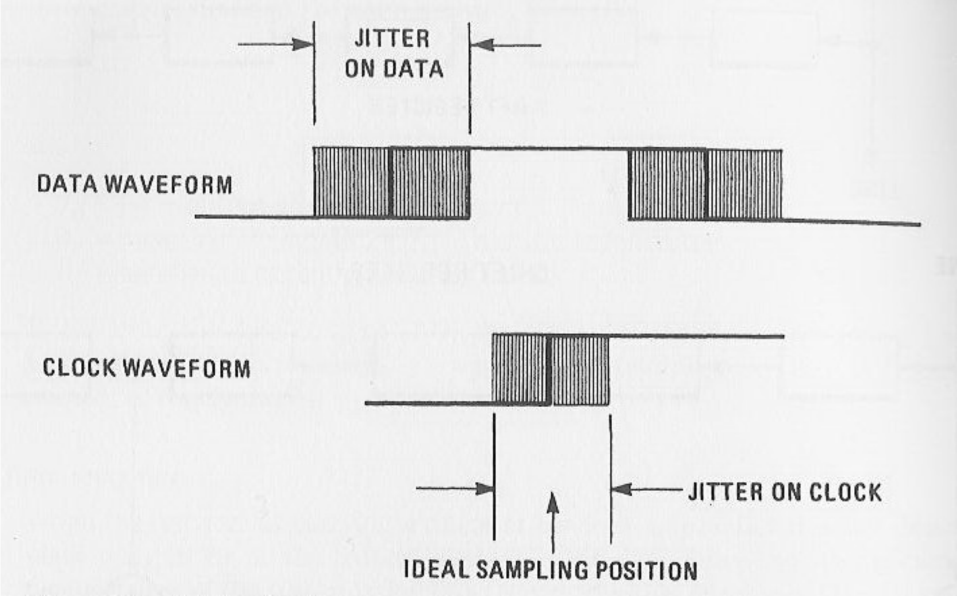

In the early days of digital communications, many links were just point-to-point, but as the network expanded in the 1970s and 80s, it was increasingly possible that multiple digital links would be connected one after another. In these larger networks, jitter began to become a problem. Jitter is caused by several mechanisms, in particular, imperfections in clock recovery circuits and the effects of buffer stores in digital multiplexers. These effects tend to build up as the digital signal transfers through multiplex equipment and multiple line sections. Jitter theory and measurements are beyond this present account, but are described in detail in Chapter 23 of the Handbook.

HP's first generation of BER testers needed to achieve low levels of jitter in both the pattern generator and receiver, particularly when the receiver incorporated clock recovery. The R&D team began to develop expertise at an early stage, and in 1981 introduced the first dedicated jitter test set, the 3785A/B. The A version catered for the international bit rates of 2, 8 and 34 Mb/s and the B version handled North American rates of 1.5, 6 and 45 Mb/s. It was thus a companion instrument to the 3781/82 described earlier, and used a similar package based on the 1700 series oscilloscopes.

The 3785 could generate a clock signal with a specific level and frequency of jitter in order to stress-test operational equipment and could make an automated sweep across the range of frequency and level to check the full equipment specification. The receiver could make very accurate measurements of inherent jitter in a clock or random data stream so the level of jitter on the data output from equipment could be checked.

These jitter specifications were laid down in international recommendations, initially by CCITT (International Telephone and Telegraph Consultative Committee), later becoming ITU-T (International Telecommunications Union – Telecom Recommendations). Engineers at HP South Queensferry were early contributors to the Study Groups that evolved these recommendations in the 1970s and 80s, specifying performance of both operational equipment and test gear[7].

Making accurate measurements of intrinsic or background jitter is very challenging as the levels permitted at higher jitter frequencies are very low. In common with other measurements, the residual error in the measurement receiver needs to be several times better than the item under test to have confidence in the accuracy. South Queensferry's jitter test sets were some of the best in the world. Occasionally, competitor's boxes gave different results, but after discussion with customers it emerged the competitor instrument had higher residual jitter.

The NMS preserved 3785A has serial number prefix 2244 indicating it was a fairly early unit manufactured in late 1982. The 3785A/B sold for around $14k and was in production from 1981 to 1990, shipping a combined total of 1750 units with revenue of $25M.

[7] The standard specifying the performance of jitter test equipment was CCITT O.171 "Timing Jitter Measuring Equipment for Digital Systems"

3764A Digital Transmission Analyzer (Unit Preserved NMS T.2010.69)

The new instrument introduced in 1984 was one of the most significant in South Queensferry's family of BER testers. The 3764A was the first to be called a Digital Transmission Analyzer since its microprocessor control allowed statistical analysis of system performance according to the CCITT error performance standards, known as G.821.

With a top rate of 170 Mb/s it was intended as a single box replacement of the 3762/3763 from seven years earlier, and apart from a few specialized applications it did just that. But it also went a lot further.

Three main versions of the 3764 targeted different markets. The standard version was a 140 Mb/s test set with clock recovery and binary clock and data interfaces, and was aimed at the mainstream trunk transmission market including digital radio. A second version (shown here) added 140 Mb/s jitter measurements, a first in the market which also rounded out HP's jitter measurements, complementing the 3785A. A third version had four internal clocks and could perform error measurements at 2, 8, 34 and 140 Mb/s in one box. The average price was around $15k.

It was an amazingly successful product with total shipments of 4700 units when it ceased production in 1996 and generating about $70m in revenue. In early 1988, the 2000th unit rolled off the production line. It was originally expected to run at 300 units per year, but by the late 1980s, the factory was getting orders for nearly 800 units per year with over 30% market share. This did reduce significantly in the early 1990s when the multi-rate version was replaced by a new product. However the specialized 140 Mb/s jitter capability continued. In some ways, the 3764 looked back to the earlier testers, being simply a PRBS pattern generator and error detector. It did not generate or monitor the multiplex frame structure described earlier.

The 3764A in the NMS collection was manufactured in mid-1992 and was a late model, being around the 4500th unit produced. It is the 140 Mb/s jitter version mentioned above. The other versions of the 3764A are preserved in private collections.

3789B DS3 Test Set and 3787B Digital Data Test Set (3787B Unit Preserved NMS T.2010.77)

While the South Queensferry BER testers were the dominant products in the European (CCITT/ITU-T) market by the early 1980s, penetration of the North American market was very much smaller. This was partly due to the factory having less contact with the market, and it was the strategy to produce a universal product and adapt it for the North American application. At that time, US telecoms was largely a monopoly run by AT&T which controlled the local telephone companies and long distance transmission. It was a harder market to break into, and in the field test application it was mainly supplied by indigenous companies who had close links with the Bell Telephone Operating Companies and Western Electric, AT&T's manufacturing operation. However, as the success of the Microwave Link Analyzer showed, it was possible to penetrate this market if the product and price were right.

Following divestiture of AT&T in 1984 and the creation of independent telephone companies or local exchange carriers (Telcos or LECs), the market seemed more open, plus several new competitors emerged in the deregulated telecom transmission and services market. HP decided to develop two new products specifically for North America (and not derived from international products), the 3789B for transmission and the 3787B for the new emerging market of digital data services.

The 3789B was a high-end tester intended primarily to connect at the DS3 or 45 Mb/s rate which was the standard interface in US transmission systems at the time. The 3787B operated at the DS1 or 1.5 Mb/s rate, also referred to as T1. Both instruments were an innovation from earlier South Queensferry BER testers as they had powerful demultiplexers and framing built into them. These allowed tributary streams to be extracted and tested as well providing monitoring and stress testing of the system framing bits. In fact the 3789B was sometimes referred to as a surveillance test set with the idea that it could scan multiple DS3 circuits through access switching to search for trouble.

Codenamed "Flossie", the 3789B was implemented mostly in discrete logic ICs, and this resulted in quite a large and heavy box. One salesman complained years later that he still had "tennis elbow" from lugging the beast around! A year or two later, the product could have used gate arrays, and been a fraction of the size. In fact timing was the Achilles heel of both products. Development started in the early 1980s, and it took four to five years to complete, far too long in a rapidly evolving deregulated market. This, coupled with some lack of market knowledge and understanding of customer needs, meant the products underwent revision during development, causing further delays.

By the time they emerged in 1987, they were a bit late in the day and probably missed some of the market. Worse still, the late arrival created ridicule from the US sales force with comments like, "*The thirty-seven eighty-WHEN?*", and others that aren't really printable. They wanted something to sell against the growing band of competitors in a rapidly expanding market, but they had to wait.

Both products were innovative in the user interface. Up until then the BER testers had used one button per function, so almost everything was on the front panel. These were the first BER testers from HP that used a menu-driven user interface. The 3789B used a 4-line alpha-numeric LCD display, and the 3787B used a small monochrome CRT. The front panels were uncluttered considering the range of measurement capability.

They were significant, being the first instruments to use a Common Firmware Platform, an idea initiated by Malcolm Rix in the 3789B team. The common firmware, called Instrument Software System (ISS), included the operating system, the man machine interface (MMI), remote control and data logging. In addition to the now reusable code, ISS made it possible to simulate the MMI for both keyboard and display. It was one of several initiatives in R&D to reuse design work to reduce duplication and improve product development efficiency. At the time, South Queensferry was ahead of many HP divisions with its Common Development Group in R&D for hardware, software and mechanical design. However, this new approach was an additional challenge for the design teams on both these products and was partly responsible for the delay in development. The new menu-driven user interface also required training of the field sales force for effective selling.

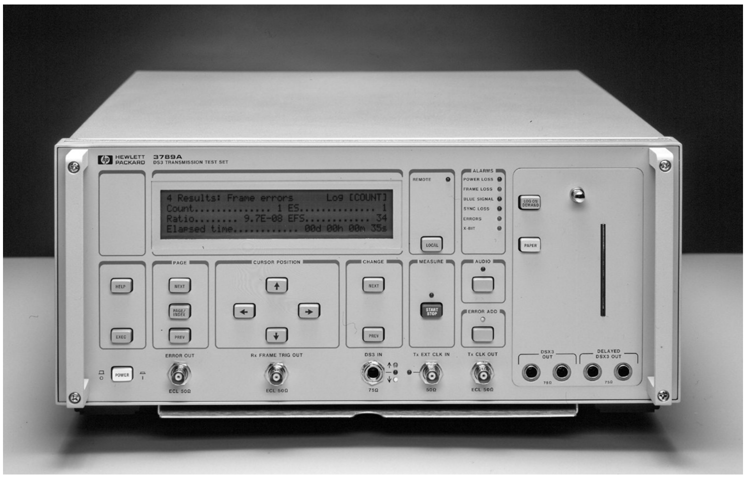

Neither product was a big success. Although the 3789B had powerful features, they were more than most Telcos wanted. The market was pretty well sewn-up by indigenous competitors: Scientific Atlanta, Tau-Tron and TTC with their T-BERD product, and later Cerjac. These products were cheaper and less than half the size.

The 3789B sold for around $12.5k and during its production life from 1987 to 1994 sold just over 1000 units producing total revenue of about $12M. This was a small share of the market, considering the market leader, Tau-Tron, shipped $11M of DS3 testers per annum in the late 80s.

The 3787B was aimed at the Digital Data Services (DDS) market in the USA. This was a new market in which low-speed data formerly carried by analogue data modems (4.8 and 9.6 kb/s etc.) could be carried over digital facilities by multiplexing them into the T1 stream at 1.5 Mb/s. This was called sub-rate multiplexing and was the particular capability of the 3787B along with DS0 (64 kb/s) and T1 testing. Given the complex set-up and choices, the CRT menu-driven display was ideal.

For some reason, the 3787B did not sell well. Possibly it was a bit expensive, but it is likely HP's North American sales force didn't have the specialised knowledge needed to sell it. This was always a disadvantage compared to the competitors like Tau-Tron and TTC, whose sole focus was the Telco market. The 3787B sold for $9k and shipped only 725 units between 1987 and 1994 with total revenue of $6.5M. Again for comparison, Tau-Tron shipped $12M of T1 testers per annum in the late 1980s.

However, the 3787B was a milestone product since it was the first unit designed as much for services as facilities test. The menu driven display set the trend for all the future BER testers developed at South Queensferry. The unit preserved in the NMS collection was manufactured in mid 1990.

3784A Digital Transmission Analyzer

In 1989, HP introduced a new product to replace several of the earlier instruments operating up to 50 Mb/s, in particular the 3780A, 3781/82A and the 3785A Jitter test Set. It capitalised on some of the circuitry developed for the abandoned 37800A project

described earlier. The 3784A benefited from this newer technology and was able to combine all this measurement capability in a single instrument the same size as the original 3780A. It had a synthesized clock source for general purpose applications. It also employed the same menu driven user interface as the 3789B DS3 Test Set, using the 4-line LCD display and navigation buttons. It allowed all the earlier instruments to be discontinued by 1991. It was quite a successful product selling 2200 units with revenue of $35M, before it was discontinued in 1996. The price was $13k for BER version and $20k for BER plus Jitter version.

A Bit of Badge Engineering

At this point we'll make a diversion from the evolution of South Queensferry's digital transmission testers to look at a bit of badge engineering the Division indulged in for a while in the late 1980s and early 1990s. No doubt this was seen as a shortcut to expanding the portfolio of transmission testers by "re-badging" other companies' products, however the experiment turned out to be a bit of a dead-end.

HP divisions, including South Queensferry, had most success with products at the top end of the market where the technology and design were more demanding and profit margins were higher. Hewlett and Packard always encouraged the idea that new products should make a "technical contribution", and not just replicate equipment already in the market. In the telecom test business, this meant the Division omitted a considerable range of low-end products which nevertheless sold in quite high volume. The idea was that perhaps this gap could be filled by re-branding, thereby avoiding most of the development cost. By this time, the Division was mainly focussed on digital transmission test sets so this was the portfolio they wished to expand. There was little point in approaching the mainstream competitors (Anritsu, Wandel und Goltermann, TTC and Tautron), so the search was for smaller start-up companies with suitable products who might be interested to benefit from HP's reputation and infrastructure. The question was whether customers would really accept these as HP products and whether the products themselves would live up to HP quality.

The first of these badge-engineered products came from an Italian company called Necsy, based in Padua near Venice. It was not a Division idea; instead it seems to have been cooked up between the HP Italy Country Manager, Roberto Favaretto, Franco Mariotti (HP Europe Manager in Geneva) and Antonella Madonna, the Necsy marketing manager. Favaretto was from Padua and went to university there and was interested in getting a joint venture going between Necsy and HP. The HP headquarters in Geneva, known in the factory simply as HPSA, always seemed to have pots of money, so they bank-rolled this “Italian job”. Robin Myles, R&D manager, recalled the first contact with Necsy (or Flitel as it was called then) took place in December 1986,

“*I remember we were driven from Milan to Padua in great haste by one of the HP Italy staff, at over 80 mph in thick fog! However, it was another year before we started serious negotiations with Necsy.*”

The proposed Necsy product was a field portable, battery operated BER tester working at the 2 Mb/s (the E1 rate) and also 704 kb/s and 64 kb/s, but without any frame structure. It was a basic instrument but complementary to some of the higher-rate testers Queensferry had produced. Its most notable feature was the robust moulded plastic case coloured bright yellow which Necsy made with their in-house plastic injection moulding facility. It looked more like something a telephone engineer would use, compared to the typical products from the Division.

Bob Thomson, one of the digital transmission R&D team, had the job of making it happen:

“*Robin Myles and I were charged with integrating this into the Queensferry product range. There was pressure from Geneva to get it done, but I remember there was a notable lack of enthusiasm from the R&D folk back at the Division, as the specification was seen as too basic. We had to figure out what changes would be required, ranging from altering the package and front panel to HP colours, to checking the operating spec was correct. I got to know a few of the Necsy staff well and they always made us feel welcome and used to fly the Union Flag alongside the Italian one outside the factory when we were there. For me it was an interesting adventure, but given the opposition to it from the home base, it did nothing to help our product line. However, I had some enjoyable trips to Padua and Venice at Geneva’s expense!*”

The 3788A Error Performance Analyzer was duly introduced in the second half of 1989 and sold for around $3k. The transformation to an HP product was effective and it had a clean, smart appearance with a horizontal LCD display for setup and measurements. The unit was manufactured for HP at Necsy’s factory in Padua.

Mature sales were forecast at 300 units per year with a market life of 5 to 10 years. Sales were a disappointment however. By 1992, the volume had dropped to 66 units and it was withdrawn in 1993 with total sales of only 330 units. It was quite well engineered, but the price was far too high for what was being offered. By the early 1990s, customers expected and needed a 2 Mb/s test with full frame generation and analysis, so it didn’t meet market needs and was quickly overtaken by new framed BER testers from South Queensferry (see the 37742A handheld 2 Mb/s tester described later).

This badge-engineering experiment probably made a loss for the Division, even with Geneva’s financial backing. However, the deal was driven more by politics than product strategy, as Finlay Mackenzie recalled.

"*Franco Mariotti was desperate to get some manufacturing activity in Italy – the other three major economies, UK, France and Germany, all had substantial HP facilities. CEO, John Young, was not keen, however he gave us the green light for a joint venture with Necsy.*"

Nothing much came of it, but perhaps it was a missed opportunity as HP became more interested in the field test market in the 1990s and Necsy might have been a useful partner – particularly with its plastic injection-mould facility!

About the same time, the Division tried a further product development experiment. This time they had another go at breaking into the North American T1 market at 1.5 Mb/s, the equivalent of the 2 Mb/s E1 sector targeted by the Necsy 3788A in Europe. Using the latest technology, it would be an early attempt at a handheld test set, and they turned to a small start-up company Taq Communications in California (closer to customers) to do the development work. Iain Milnes, a design engineer who had worked at South Queensferry, transferred to HP in the USA in the mid 1980s and later set up Taq[8] in the late 1980s.

Milnes came up with an innovative product which used quite a large LCD display panel and six navigation keys to scroll through set-up parameters and measurement results. Unlike the Necsy box, it was a fully framed BER test set, and could auto-configure to the received signal. For its time it looked an elegant design, but Milnes had cash-flow problems and needed advances from the Division to complete the task. It was introduced in 1991 and sold for about $4k, however for some reason it appeared to be a poor seller initially, achieving only 97 units in 1992 and 160 in 1993, a small share of the T1 tester market. It was up against some stiff competition from a similar product supplied by a US company, Sunrise Telecom. In 1994, responsibility for the product transferred to the new Cerjac Operation in the USA (an acquisition discussed shortly), so subsequent sales volume is unknown. It did stay in the HP catalogue until 2000, so must have done reasonable business. The product, the 37741A, is seen here being operated by Alistair Calder.

There was quite a change for the Division in 1990 when the General Manager and one of the "founding fathers" of HP at South Queensferry, Finlay Mackenzie, handed over the reins to an American, Chuck Acken. Finlay always had a strong interest in product strategy even as general manager, whereas Chuck was more focussed on sales and marketing. In particular he wanted to improve relations with the field sales force, which had been a bit strained during the 1980s. As HP morphed into a computer company, the field reorganisation that followed hadn't helped the telecom specialisation that was needed to sell against dedicated competition. South Queensferry was also perceived in some

[8] Some at South Queensferry remember Milnes as a clever guy, but a bit of a maverick and a "driven man" who didn't fit comfortably into HP corporate culture. His first venture, Taq, got into financial difficulties in 1994 and was "subsumed" into a new venture called Zarak which had great success in the late 1990s with a modular product called "Abacus" which provided load testing for telecom switches. In late 2000, at the peak of the dotcom boom, Zarak was acquired by Spirent/Bowthorpe for $400M in shares, a major portion going to Milnes. A true entrepreneur, a year later in 2001 he got into IP telephony and set up another company, Zultys, specialising in IP PABX. It was less successful and was sold at auction in 2006 for $1M after filing for Chapter 11 bankruptcy.

quarters as not being very responsive and producing over-complicated products late to market. Chuck Acken tried a number of initiatives in the 1990s (described in Volume 1 Chapter 10) to woo the field, and also decided to give them some extra products – quickly – by badge engineering.

The team identified two manufacturers in the USA and one in Europe who looked willing to get involved in an OEM[9] deal. In the USA they found a low-cost DS3 test set from a company called T-Com, and an early SONET test set from a company called Antel. These were repackaged in the current portable instrument housing used at South Queensferry and given revamped front panels to make them look like in-house products. The T-Com DS3 tester became the 37743A and the Antel box the 37744A. The 37743A was priced at $7.5k, so was considerably cheaper and smaller than the Queensferry designed 3789B described earlier. It also had mux and demux capability. Introduced 1991, it didn't sell well only achieving 49 units in 1992 and 24 in 1993 – a complete failure in fact. Either customers didn't like what they saw, or the field felt uncomfortable selling it against the entrenched competition. Some commented that the quality wasn't up to HP standards. It was a similar story with the 37744A SONET tester (shown here) which was only on the market for a couple of years, despite an attractive price of $10k. A disappointing outcome, which failed to get any market penetration.

The European company in the OEM deal was ICT of Spain. They supplied a pair of units for generating and analyzing the European frame structure from 2 Mb/s to 140 Mb/s. After cosmetic revamping, these became the 37729A Frame Generator and 37730A Frame Analyzer. Although the idea was fine, the ICT products were quite bulky and heavy (probably old technology), and were expensive at $22k for the generator and $18k for the analyzer. Introduced in late 1991, the forecast on the 37730A Frame Analyzer was an ambitious 20 units/month! Neither product did much business during the two years they were in the market, the 37729A notching up sales of 40 units and the 37730A 150 units. A complete waste of time and money, and one has to remember that each of these badge-engineered instruments required someone to produce sales literature, an operating manual, and devise a plan for post-sales support of a product not manufactured by HP.

The third product from ICT was a 2 Mb/s handheld tester, similar in concept to the 37741A Handheld T1 Tester described earlier. Compared to the T1 unit, it was a bit chunkier and had a more conventional front panel with a membrane keyboard and LED status indicators. There was also a small four-line LCD panel for results. The unit was very little modified apart from adding the HP logo and product number, which was the 37742A. It was far more capable than the earlier 3788A product from Necsy, having framing and mux/demux to 64 kb/s channels. It came on the market in early 1992 selling at around $5k list price and with forecast demand of 24 units per month.

This time the badge-engineering team had picked a winner! It quickly grew to 30 units per month and more by the mid-1990s. Between 1992 and 1997 it sold 2300 units and

[9] Original Equipment Manufacturer, a term used in the industry to denote a supplier of assemblies or complete products that were re-branded to be sold under a different label.

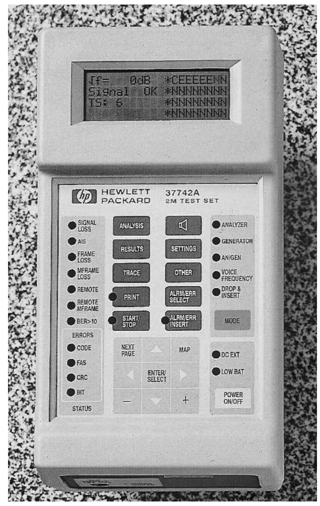

continued in the portfolio until around 1999 when it was replaced by the Queensferry-designed ProBER2 product, E7580A (described in Chapter 10). South Queensferry found it far easier to penetrate the market outside North America, and along with the 37722/32A testers, developed as part of the Hornet programme and described in the next chapter, the Division had a substantial share of the 2 Mb/s E1 tester market, shipping over 100 units per month in total.

A small engineering team headed by Robert Duncan was set up to negotiate these OEM contracts with the third party companies, and to do the adaptation necessary to make the products look like HP kit. The alterations were in most instances skin-deep.

Quite a lot of antipathy developed between this team and some in the R&D department who despised the whole re-branding experiment and saw some of the products as "dross". In turn this was viewed as "sour grapes" and a case of the "Not Invented Here" (NIH) syndrome. NIH apart, the South Queensferry R&D team could have designed this stuff standing on their heads, but the fact was they weren't interested in developing such basic products, and the Division didn't have the resources to do everything. They also rather missed the point. Chuck Acken wasn't looking for technologically smart products; he wanted some simple things that customers had proved they wanted to buy and field sales wanted to sell! Unfortunately, it turned out to be a lot harder than he thought.

He made one further move, the most ambitious so far, to penetrate the US telephone operator market by acquiring Cerjac in July 1993. Cerjac, a private company based in Westford near Boston Massachusetts, had been in business several years and had developed some products for T-carrier and SONET testing that had been successful competing with mainstream suppliers in North America. In this case the products retained their identity and were marketed as HP Cerjac[10]. Chuck's master plan was to transfer responsibility for North American instruments from South Queensferry to the new Cerjac Telecom Operation to get the business closer to the customers. This caused more unrest in R&D.

In the end, the stand-off became irrelevant. As we have seen, most of the badge-engineered products were a disappointment and while this was going on, R&D was developing a new product platform for the digital transmission analyzer family which rather eclipsed all these third-party activities. Codenamed "Hornet", it helped to power the South Queensferry Division to a point in the late 1990s where it became probably the global market leader in digital transmission test, as we shall see in the next chapter.

Acknowledgements

Thanks to the following South Queensferry employees who helped with this chapter: Bob Thomson, Tom Crawford, Ivan Young, Finlay Mackenzie, Robin Myles, Andrew Wilson and Peter Scott

---

[10] Two key products were the HP Cerjac 156 SONET and T-carrier test set and the stripped-down version, the HP Cerjac 156MTS, and were selling about 20 to 25 units/month.

10

Chapter Ten

Digital Transmission Analyzers 1990 – 2005

1990 marked a significant watershed for South Queensferry's family of Digital Transmission Analyzers. Some of the changes were technical and technological, and others were market driven. This chapter tells the story of how the Division reacted to these changes and generated a great deal of new business.

The New Wave

In the late 1980s, the telecom industry was migrating from the first generation Plesiochronous Digital Hierarchy (PDH) to the new synchronous multiplexing standard. This was known as Synchronous Digital Hierarchy (SDH) internationally, and as SONET (Synchronous Optical Network) in North America. At this time, the technology of fibre-optic transmission was advancing rapidly with the introduction of single-mode fibre and narrow-band laser transmitters. This meant bit rates moved quickly into the Gigabit per second range, whereas most PDH systems had operated at 140 Mb/s and a few at 565 Mb/s. The market required a new generation of test equipment that could operate at higher bit rates, interface with optics and also generate and monitor the many complex overhead signals in the new synchronous frame structure. However, these instruments also had to be backward compatible with the older PDH installed base.

Up to this time, all the classic products described in the last chapter used conventional printed circuit board construction and with standard integrated circuits and other discrete components inserted through holes in the board and soldered. Production methods had moved from manual component insertion to the use of computer-controlled auto-insertion machines, but the end result was similar. All products were shipped with a comprehensive service manual which included circuit diagrams, descriptions and component lists, so the product could be repaired down to component level in a service office.

From 1990 onwards, all new instruments used Surface Mount Technology (SMT) for some or all the circuit boards, in line with common practice in the electronics industry. In SMT, much smaller components (for example chip capacitors and resistors and miniaturised integrated circuits) were positioned on the surface of the board by a “pick-and-place” machine and held in place by a small blob of adhesive, the connecting pads on the board having previously been coated with a solder paste. The loaded board then went through a re-flow soldering oven which raised the temperature to the point where the paste melted and soldered the components to the board.

SMT had several advantages. It was ideal for automated production of loaded boards through “pick-and-place” machines and the end result was an assembly which was a fraction of the size of the old-fashioned through-hole soldered PCB. The smaller size of SMT also meant it potentially had better electrical performance because of lower stray capacitance and shorter propagation delay – ideal for the higher bit rates now required.

Various custom integrated circuits which emerged in the 1980s were used extensively in the new products. These were the Field Programmable Gate Array or FPGA with companies such as Xilinx and Altera being key suppliers in the early 1990s and also the Application Specific Integrated Circuit or ASIC. Both these custom ICs allowed much more complicated logic functions and communications protocols to be implemented in a small number of miniaturised packages. Together with SMT, this technology greatly increased the functionality of the new generation of digital transmission analyzers developed after 1990. Here is a photo of part of a circuit board in the 37724A SDH/PDH Test Set from 1992 (described later) featuring a number of Xilinx FPGAs and some ASICs. It was one of the first products to use this new technology.

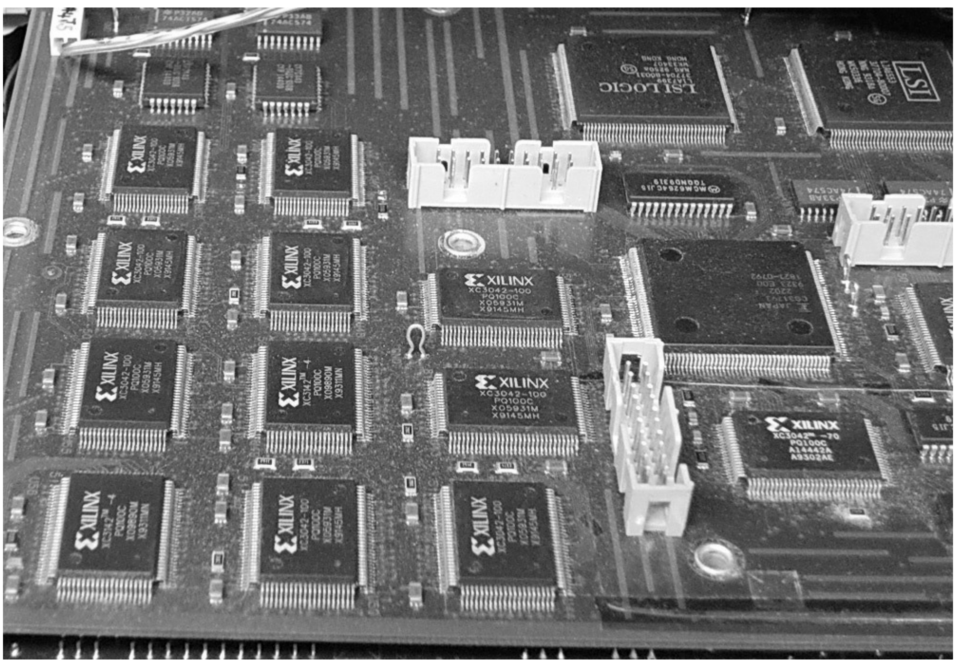

The downside was that the new products could no longer be repaired in a service office to component level, as SMT required specialised facilities. South Queensferry sub-contracted production of all the SMT assemblies, although a department was set up at the factory with the necessary equipment to repair boards on site as part of the production process. It also meant prototyping in R&D was more difficult, so more use was made of computer aided design in the hope that the externally manufactured prototype boards would work first time. Similarly, all design work for FPGAs and ASICs needed to be done on a computer work station. Instruments could only be repaired by board or module exchange, so the products were no longer shipped with circuit diagrams in the service manual. Everything about the new products was markedly different from the classic products designed in the previous 20 years.

This advanced technology and new market landscape was the backdrop to a new instrument platform which was evolved in the late 1980s by South Queensferry's R&D engineers. This was called "Hornet". In fact it was a hallmark of this new era in that the engineering team developed a penchant for in-house "codenames" in the 1990s, as we'll see a little later.

The "Hornet" Family

The final group of classic products, 3789B DS3 Test Set, 3787B Digital Data Test Set and 3784A Digital Transmission Analyzer described in the last chapter, confirmed the need to move to menu-driven user interface as measurement complexity increased. Even the 3789B and 3784A were a bit awkward to use and the legibility and visibility of their LCD displays was poor compared to the earlier LED displays. The future seemed to be the more flexible CRT display of the 3787B.

After the prolonged development times for the 3787B and 3789B, division management were also very aware of the business penalties of failing to reduce time to market. Although these two products came from the same group in R&D, they didn't appear to have much in common. Tom Crawford pushed to develop a common platform for future products that would reuse the man-machine interface, mechanical design and where possible items like power supplies and other common circuitry. The instrument case was a portable package borrowed from HP's 8590 series of spectrum analyzers. As Tom Crawford recalled, the objective was to improve R&D productivity through shared ways of working. The Common Development Group was established to encourage project managers to engage in the success of the product family as well as their own product or module. A project manager could also be asked to head the development of common-technology to be reused across a group of products. So the project managers became necessarily dependent on mutual help and sharing.

An early contribution to the common platform idea came from Malcolm Rix and Nick Race who invented and developed a common software platform, optimised for real-time instrument use. The Division also embarked on a new venture called New Product Introduction (NPI), later Fast Cycle Time (FCT), that looked at all aspects of the product development process, aiming to reduce timescales. The objective was to get overall development time from definition to product launch down to two years, and less if possible. This is described in more detail in Volume 1 Chapter 7.

This new platform, developed in the late 1980s, was codenamed "Hornet". In its various forms, it was used throughout the 1990s for a whole new family of digital transmission analyzers and was probably overall the most successful development ever undertaken at South Queensferry.

37721A Digital Transmission Analyzer

The first member of this new family emerged in 1990. This was the 37721A Digital Transmission Analyzer. The R&D team was on familiar ground with this product as it was a revamp of the multi-rate version of the 3764A unit introduced in 1984, and the market leader for PDH transmission test up to 140 Mb/s. The new product used newer IC technology and SMT boards, so was more compact, and for the first time included analysis of the frame structure transmitted by operational equipment. This gave it some in-service measurements in addition to the traditional pattern generator and error detector testing which was always done out-of-service.

The most striking feature of the new product was the electroluminescent display and the row of “softkeys” along the bottom which altered according to the application or parameter selected. The monochrome yellow display was bright and highly legible and could switch to large bold characters to display key measurements, a big improvement on the earlier products. Two navigation buttons allowed different fields on the display to be selected for setting parameters using the appropriately labelled softkeys[1]. The display could also present measurement results graphically. This combination of keys and display was the blueprint for the new series of transmission testers that followed. Here is an example showing how the highlighted data field in the display sets a specific group of soft-keys at the bottom of the screen:

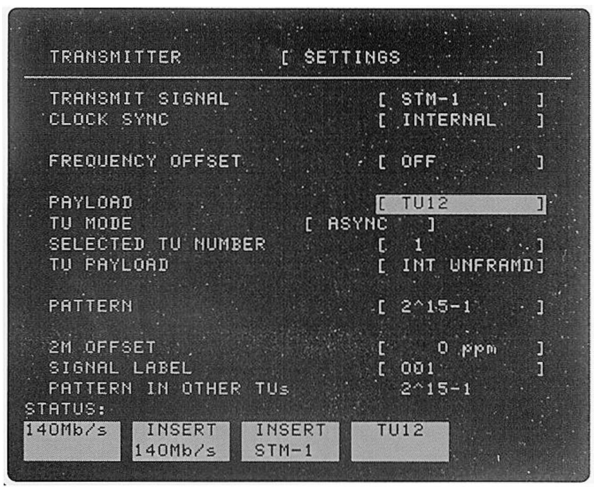

In late 1989, the project leader, John McElroy, visited Telecom Australia with two 37721A prototypes. The customer was very impressed with the new product, particularly the display and man-machine interface (MMI), with considerable praise for its implementation and ease of use. This successful visit led to the promise of an order of 100 to 150 units the following year. The 37721A sold for around $12k and 3500 units were shipped between 1990 and 1999 with total revenue of $45M.

[1] Tom Crawford recalled they made two experimental versions of the front panel, one with UP/DOWN keys and one with an RPG knob for navigation, and tested them using a group drawn from around the factory. “The RPG lost … I felt it was the best … but there we are!”

37701A/B and 37711A T1 Test Sets (Unit Preserved)

Around the same time, South Queensferry decided to have another go at the North American market and developed another T1 tester for 1.5 Mb/s and below. As mentioned earlier, the Division's somewhat unsuccessful forays into the North American market were caused partly by strong local competition and also by incomplete market knowledge. The digital transmission systems there were somewhat different to those in Europe. Apart from the obvious differences in standard bit rates, the North American T1 systems were more pattern-dependant due to the line codes used (Alternate Mark Inversion and B8ZS). Whereas the European systems were designed to be completely bit-sequence independent (i.e. transparent) and could be reliably tested with any arbitrary bit-pattern, the T1 systems required several specific patterns for stress-testing which complied with the bit-sequence restrictions. For example, one such pattern was called Quasi Random Signal Source (QRSS) which was a 20-stage PRBS limited to a maximum run of only 14 zeros.

The new product to satisfy all these needs was the 37701A T1 Test Set. It used the same Hornet front panel and MMI, but was very compact having an overall depth of about 8.5" (210mm) and weighed around 4.5 kg. It cost around $6k. A second version, the 37711A, included a basic datacom test capability in the lid (15901A with RS232, V.35 etc.). This sold for around $8k. In 1993 they were superseded by the 37701B T1/Datacom Test Set and the 37702A Digital Data Tester which replaced the earlier 3787B. The prices were $9k and $12k respectively. Between 1990 and 2000, these T1 variants sold around 3000 units in total, bringing in around $30M. A reasonable result, but it still represented a small share of the North American T1 test market. It remained very difficult to penetrate the entrenched local competition.

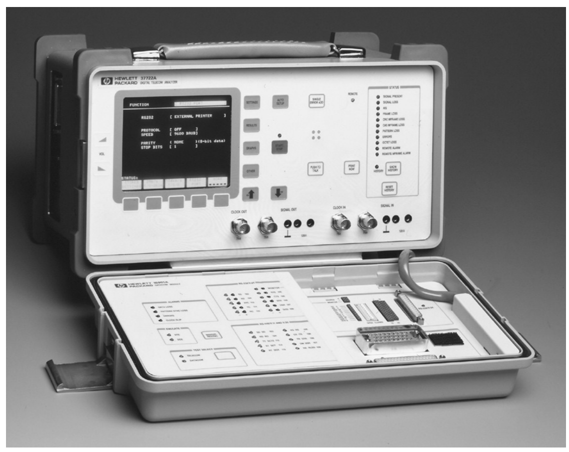

37722A/32A Digital Telecom Analyzer (Unit Preserved)

Much more successful was the European version of the same product, the 37722A E1 Tester and the 37732A with the added datacom measurements in the lid. It operated up to 8 Mb/s but was mainly intended for 2 Mb/s and below, targeting the digital private line market such as BT's Megastream and Kilostream services. Both instruments were introduced in 1992.

The 37732A E1/Datacom Tester (shown here) was the big seller, achieving 5000 units by the end of run in 2000. Together, these two products brought in over $50M. Along with the badge-engineered 2 Mb/s handheld product from ICT (37742A) described in the previous chapter, the Division did good business in the E1 tester market throughout much of the 1990s, as shown in this graph.

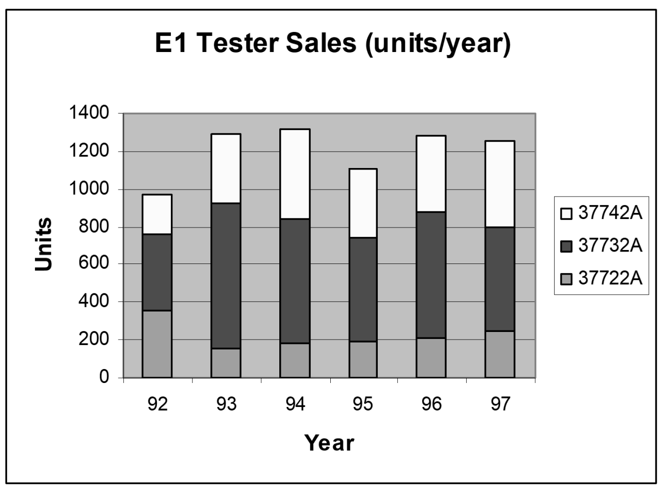

SONET/SDH

The first generation of digital telecom systems in the 1970s and 1980s were built using the Plesiochronous Digital Hierarchy or PDH. Plesiochronous means almost synchronous, that is the transmission bit rates are very close to the nominal rate but are not required to be exact. This does not apply to the primary rate of 1.5 or 2 Mb/s which is also the rate used in digital telephone switches. The switches and therefore the 2 Mb/s equipment is all synchronised to a network master clock. For the hierarchy rates above this (8, 34 and 140 Mb/s), PDH applies.

The original idea was that PDH would make for a more robust network since the multiplexing equipment could operate autonomously and would automatically compensate for variations in transmission delay, frequency drift and so on. It did this by synchronising the incoming data stream to its own internal clock by adding or subtracting specific bits in the stream using a buffer store. These bits are called the “stuffing bits” and the process is known as justification. It works well, however the downside comes if you want to extract say a particular 2 Mb/s stream multiplexed in a 140 Mb/s transmission. To do this, you first need to systematically reverse the justification process stage by stage. In the early days this didn’t matter much, but as networks became more complex, operators wanted to manage and allocate bandwidth more flexibly.

To answer this need, a new standard was evolved in the mid 1980s called synchronous transmission in which all multiplex equipment at all hierarchy levels was synchronised to a master clock. This meant there was no need for the justification process, so any tributary stream can be identified and extracted simply by locking onto framing words and other overhead bits. Two parallel standards evolved that are broadly similar. In North America it is known as SONET (Synchronous Optical Network) and within ITU-T as SDH (Synchronous Digital Hierarchy). Since the new generation systems would invariably operate over fibre cable, the standards define not only the synchronous frame structure but also the optical transport. A full description of SONET/SDH background is given in Chapter 13 of the Handbook[2].

As SONET and SDH systems developed in the 1990s, a number of further enhancements were added to the standards to handle different types of telecom traffic. The synchronous frame structure comprised two parts: the overhead bytes used by the transmission system for management and identification, and the payload which was the revenue-earning customer traffic. A popular analogy was the articulated truck: the transport overhead bytes that made the thing go, and the container on the back which was the payload. The payload area of the frame could be subdivided in different ways (called payload mapping) to suit different traffic. In the early days it was mostly telephony and a common mapping was to carry lower-rate PDH traffic. Later, as broadband data traffic predominated, new mappings were standardised for taking Asynchronous Transfer Mode (ATM) cell traffic mentioned later, and then as Internet Protocol (IP) traffic came to dominate, there were mappings for putting IP packets into the payload. In the USA this was known as Packet over SONET (POS). As South Queensferry’s products developed in the 1990s, these various mappings were available as part of the new SONET/SDH test sets, and Internet router manufacturers like Cisco became major customers for the instruments.

---

[2] “*Communications Network Test & Measurement Handbook*” by Coombs and Coombs, McGraw-Hill (1998) ISBN 0-07-012617-8.

Finlay Mackenzie, Division Manager at the time, recalls some interesting challenges and opportunities resulting from the move to SONET/SDH. The demands of the frame structure described above with the ability to multiplex and de-multiplex, meant the design team needed to develop some custom Integrated Circuits for a practical implementation. This would be very expensive, however one of the design engineers, Gordon Rhind, visited Alcatel in the USA where that company was developing ICs for operational equipment. They offered to supply suitable chips to South Queensferry if the division paid $600k and sent an engineer to their facility in Rayleigh to include the specific test capability.

Finlay asked his corporate managers if they could assist with this up-front investment. None was forthcoming, and although very focussed on the photonics market in the late 1980s, the corporate folks didn't seem to appreciate the significance of SONET. "Surely, you're not getting into something else?!" Maybe they thought it was poems written by Shakespeare that the "Brits" were getting into. Fortunately, Finlay negotiated 50% of the investment from a development fund in HP Geneva (HPSA) and the rest came from the Division. The deal with Alcatel was signed in May 1989. South Queensferry negotiated a second technology transfer agreement with Alcatel in France for the supply of SDH ASIC[3] chip sets and technical support. This agreement was signed in June 1992 and the contract cost 520k Euros.

37704A SONET Test Set and 37724A SDH/PDH Test Set (Unit Preserved NMS T.2013.62)

South Queensferry's first stand-alone SONET and SDH testers were introduced in 1992. Based on the Hornet platform and user interface, these portable instruments combined some PDH testing with SONET or SDH up to 622 Mb/s, depending on the optical interface plug-in fitted. In the case of the 37724A SDH/PDH Test Set it could also make PDH measurements with framed or unframed test signals at 2, 34 and 140 Mb/s, presumably reusing the circuitry developed for the 37721A PDH tester introduced a year or two earlier. Because of the unified standards, both the 37704A and 37724A could use the same optical interface modules (the 37772A for 52/155 Mb/s and 37776A for 155/622 Mb/s).

While the 37704A had fairly modest sales of 300 units in North America, the 37724A SDH tester, for the European/International market sold well, shipping 1500 units between 1992 and 1998 with revenues of around $75M – a very good result for an emerging market. It sold well into both the network operators and equipment manufacturers. Peter Scott, the project leader, recalled that the instrument even sold to Deutsche Telekom, normally the preserve of South Queensferry's German competitor Wandel & Goltermann. The preserved unit is a prototype, c.1991.

[3] Application Specific Integrated Circuits

75000 Series 90 System

About the same time, HP introduced a VXI system for SONET/SDH and ATM test. The industry-standard VXI mainframe was a 19" rack card cage including a controller and all required power supplies and cooling. Into this could be plugged any configuration of VXI compatible modules from HP or other vendors (at least that was the intention). HP produced a wide range of VXI modules covering basic measurements. South Queensferry collaborated with HP's Australian Telecom Operation in Melbourne to develop modules for a SONET/SDH test system working up to 622 Mb/s initially. Australia's particular interest was the first generation of broadband communications based on Asynchronous Transfer Mode (ATM) which transmitted the broadband data stream in predefined cells with headers. In the late 1980s and early 1990s, South Queensferry was involved with Australia in an EC funded R&D project called RACE to develop broadband systems. One of the Division R&D engineers, Gregan Crawford, worked closely with this initiative and built up a lot of knowledge of ATM technology and the key players in the market. The ATM test contribution from HP and other companies was called "Parasol". See Volume 1 Chapter 9, pages 279 to 280.

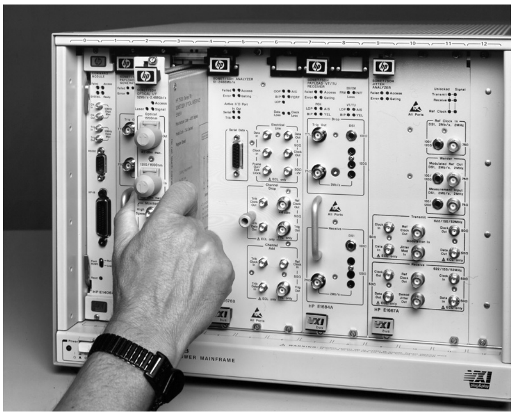

The 75000 Series 90 VXI system had the configuration flexibility to cope with many different applications. It operated with a PC (typically an HP Vectra) and software to provide the user interface and control, and was often sold as a bundled system with the VXI mainframe. The first offering was for ATM test in 1991/92, and soon a module was available for working up to 2.5 Gb/s SONET/SDH at optical interfaces (E1675A), probably the industry's first STM-16/OC-48 measurement solution. Bob Thomson, the project leader responsible for this module, recalled it was one of the most satisfying periods in his career, and enjoyable working with the Australian operation.

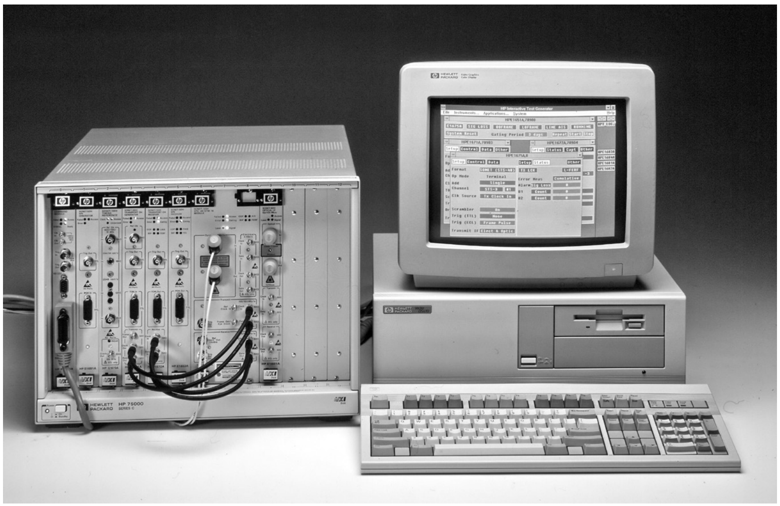

During the 1990s many other modules were introduced from Queensferry and Australia which broadened the application for measurements such as jitter and PDH. A fairly complete listing of the many VXI modules in this system has been compiled in the spreadsheet catalogue in Appendix 1, with pricing and unit volumes where known. With incomplete sales information and varying configurations, it is difficult to know how much business was done with this system. It was in production from 1991/2 until 2000 when all items were withdrawn. The complete bundled systems typically sold for $100k to $200k, and it is likely total revenue was around $100M.

37778A STM-16/OC-48 Transmission Test Set

Capitalising on the 2.5 Gb/s VXI modules, South Queensferry decided to package these separately in a Hornet style box. The modules plugged in at the rear which is where the electrical and optical signal connections were made. The power supply, cooling and control hardware was all between the front panel and the rear of the VXI plug-ins, making it one of the longest products produced at the Division, around two feet (600mm) in length! This was the 37778A STM-16/OC-48 Transmission Test Set with capability of transmit/receive functional test and also jitter testing at 2.5 Gb/s (another first). Despite its ungainly size, technical merit ensured sales of around 400 units between the mid 1990s and 1999, when it was discontinued in favour of the new 37718A described later. It sold for around $110k, giving revenue of $45M. Ironically, more 2.5 Gb/s VXI modules were sold in this cobbled-together package than in the VXI mainframe! Another win for the Hornet recipe.

75000 Series 95 and SpectralBER

In 1994 a further iteration of the VXI system was introduced specifically for functional test of multiplexers and switches in production, called 75000 Series 95. This added modules for fan-out and switching of digital and optical signals. The bundled system with controller was marketed as TS-2000. It was not a large seller, and the system and modules were discontinued in 1998.

A final variation on the system came in the late 1990s when new modules were added for testing Wavelength Division Multiplexed (WDM) systems operating up to 10 Gb/s. This was marketed as SpectralBER and was introduced in 1999/2000. A number of the high-speed optical modules were badge-engineered VXI units from a Japanese competitor, ANDO, which South Queensferry collaborated with. The Japanese were keen on building relationships and frequent exchange visits were organised between South Queensferry staff and their opposite numbers in Japan. As part of the "ritual", gifts of Japanese Daruma[4] Dolls appeared on the desks of various managers and engineers, the size relating to hierarchical position. The idea was that one eye was coloured in at the start of the project for good luck and the other coloured in when the task was completed successfully. The joint venture and the product were launched at Telecom '99 tradeshow in Geneva, with displays of the instrument on both company stands.

Initially the SpectralBER system sold quite well, just before the "dotcom" bubble burst, but for some reason the joint venture was discontinued by the end of 2001, possibly because the bottom fell out of the market or maybe because ANDO was taken over by Yokogawa in 2002. Perhaps the Daruma Dolls never had the other eye coloured in? In any case, the original Series 90 VXI system ceased production in 2000 since its capability was by then largely replaced by the highly successful Modular Hornet family, discussed next.

---

[4] Darumas are supposed to symbolise good luck and perseverance in Japan. The ones at South Queensferry were really just a head, mostly red and gold in colour with a cartoon kind of face. The accompanying instructions stated the official name as "Fuku Daruma" (good-luck Daruma), which caused some hilarity at the Division!

Modular Hornet

The first Hornet products, described above and introduced in the early 1990s, were conventional in assembly. Each product had a more or less fixed configuration (except for the two optical plug-ins in the SONET/SDH testers) and the signal connectors were on the front panel, as they had been in the past. Space was a bit of problem on the 37724A SDH/PDH Test Set, and several PDH connectors had to be placed on the rear panel, so it was quite likely that the user would have to make connections to both front and back of the instrument simultaneously when making tests.

Another issue was the rapidly evolving requirements in the market as it moved from PDH to SDH, and the question of test equipment becoming obsolete. The 75000 Series 90 VXI system solved that problem as it could be reconfigured and new modules could be added for changing test requirements. This flexibility was not available on the original Hornet boxes.

The solution to all these issues came with the development of Modular Hornet. This retained the control features and man-machine interface (MMI) of Hornet, but now the measurement hardware was built into a number of factory-fitted modules that plugged vertically into the side of the instrument. Tom Crawford commented that the design had to ensure future compatibility. He and Barry Smith devised a passive backplane (i.e. pure interconnection without electronics) – passive because the cost and space penalty in each plug-in with an active approach was too high for a field portable product.

Each module had the necessary connectors for that particular function, so there was never a problem accommodating the connectors on the instrument front or rear panel. The front panel could be devoted entirely to the user interface without the clutter of cables. All connections were to the side of the instrument at the modules, which was fine for stand-alone use, while the generic front panel with soft keys could adapt as different software was loaded on the instrument depending on the configuration.

Customers could return their instruments to the factory for upgrading, thus protecting investment. Perhaps few customers actually sent their instruments back for upgrading but it was reassuring to know they could. The greatest benefit of Modular Hornet was probably for the production system itself, which could relatively easily configure a wide range of customer specifications by assembling the appropriate plug-in modules and loading the matching software.

It was a brilliant design solution that provided better value for customers and better productivity and flexibility in the factory. However, as John Wotherspoon from product marketing recalled, in his view the future was somewhat uncertain:

"*I can remember presenting the Modular Hornet business plan to Tom White, marketing manager at the time. Tom was hell-bent on moving everyone to system test solutions, stating it would be the last portable instrument product that he would fund. Fortunately its success helped finance other portable products later in the decade!*"

37714A and 37717A/B/C PDH/SDH/ATM Tester

The first modular product, introduced in summer 1993, was the 37714A PDH/SDH/ATM tester. It had optional modules for SDH and ATM. It was advertised at the time as "*The PDH tester with a future*", as it could be upgraded with the new modules. It is shown here in the photo used for the advert. It was a good idea, however it soon became obvious that the 37714 had insufficient module space to cater for all the possible future requirements including jitter measurements. Its big brother, the 37717A introduced around the same time was what the market wanted as its longer enclosure could handle a greater number of modules. Both units could operate up to 622 Mb/s (STM-4) with the appropriate module fitted. While the 37714A only shipped 150 units, the 37717A did nearly 1400. In June 1995 it was superseded by the 37717B and around the same time modules for jitter measurements up to 622 Mb/s (STM-4) were added.

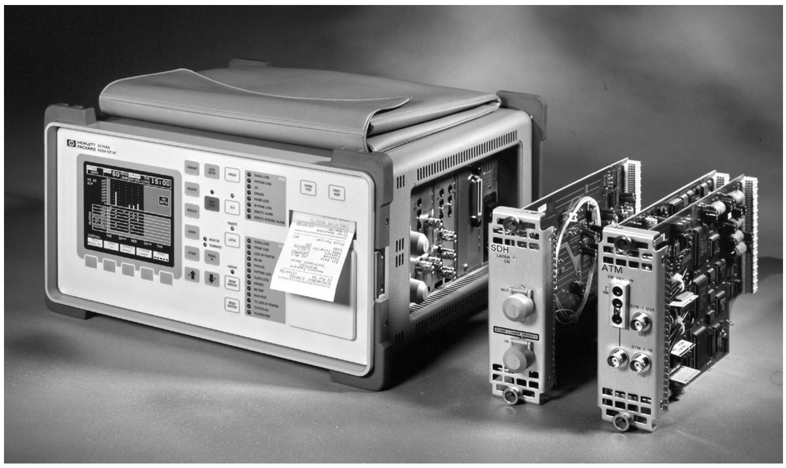

In summer 1996, South Queensferry announced it had won a substantial contract to supply the SDH test equipment for a mammoth optical submarine cable system called FLAG, a $1.5B joint project between AT&T Submarine Cables in New Jersey and KDD in Japan. The system was in eight sections, crossing the Mediterranean, Red Seas, Atlantic, Indian and Pacific Oceans, running at 5.3Gb/s over two fibre pairs. Testing was done at the 155 Mb/s STM-1 interface, and HP supplied 50 test sets, a mixture of 37714As and 37717Bs when jitter measurements were needed.

Along with the test sets, South Queensferry also supplied its recently introduced Distributed Network Analyzer (DNA) Software[5], E4540A. This was also called Virtual Remote, since the Windows-type application replicated the remote instrument's front panel at the user's PC and buttons could be pressed using the mouse. Certainly, the remote terminals on the FLAG system were an ideal application. It was introduced in 1994 and the software package sold for around $7k.

In late 1995, the 37717C was introduced. This was a significant enhancement as the electroluminescent display was replaced by a larger full colour back-lit LCD panel. This new display allowed four windows to be displayed simultaneously. As with all these Modular Hornet instruments, the selling price could vary considerably depending on the option configuration. For example, when introduced, the 37717C had a typical selling price of $35k, but could cost up to $87k with the most expensive optical and jitter measurement options fitted.

In 1998, the membrane keyboard that had been used from the start of the Hornet family, was replaced by a rubber keypad with a more "tactile" feel. The new version was re-branded as "OmniBER 717", and could optionally operate to 2.5 Gb/s (STM-16, OC-48).

[5] Described in the HP Journal, October 1994, p75-82
http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/1994-10.pdf

The 37717A/B/C was a highly successful development with total shipments of over 7000 units and revenue of around $300M. By 1996, the combined sales were running at 100 units per month, with the "C" model gradually replacing the earlier versions. The 37717C (OmniBER 717) continued in manufacture until 2000, when production rolled over to the OmniBER 718.

OmniBER 718 Communications Performance Analyzer (Unit Preserved NMS T.2010.72)

This was the final evolution of the Modular Hornet platform introduced in 1998. It looked similar to the 37717C but with some enhancements. The "OmniBER 718", codenamed "Firefly", came in three versions. The 37718A was the top of the line model with full capability up to 2.5 Gb/s, and the biggest seller. The photo below shows this model along with a typical module configuration. The list price on this was $100k for the average configuration, but could go to $180k with the most expensive options including fancy optical interfaces, although these prices were typically discounted around 20% on many deals.

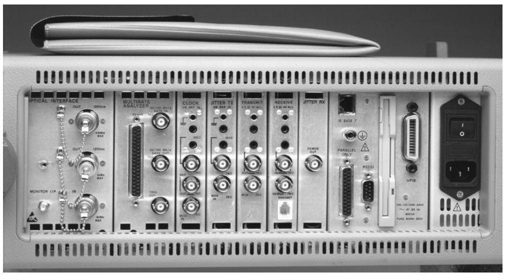

The volume on the 37718A (the 2.5 Gb/s model) was a staggering 1437 units in the boom year of 2000, when the total orders for Modular Hornet products reached 3150 units with revenues $216M. The B version went to 622 Mb/s (codename "Locust" presumably because it was lower-cost) and C version to 155 Mb/s.

Further enhancements in 2000 added additional payload mappings such as Packet over SONET (POS), ATM and an "Advanced Payload Engine". This version was codenamed "Dragonfly". The OmniBER 718 products were highly successful, and a major contributor to the enormous boom in sales experienced by the Division in 1999 and 2000 during the "dotcom" bubble. For a year or two, the South Queensferry operation was the most profitable in the HP Corporation.

Given the success of the OmniBER 718, a number of variants were developed. The OmniBER 719 was a version specially targeted at the North American SONET market and included Packet over SONET (POS) mapping described earlier. Between 1999 and 2002, it sold 590 units. Another couple of variants, the 37720 and 37725 introduced in 2001, were targeted at some production markets previously covered by the 75000 Series 90 VXI system, however by the time they were launched the market had gone cold and they never did very much.

The success of the "OmniBER 718" family can best be summarised in its sales figures.

Model	Units Sold	Revenue
37718A (2.5G)	3800	$380M
37718B (622M)	1220	$52M
37718C (155M)	1370	$57M
37719A/B/C	590	$47M
Totals	**7000**	**$530M**

As the bubble burst in 2001, sales dropped 40% compared to the previous year. A report in 2001 noted that Cisco was still the number one global customer for OmniBER and that 50% of the sales were in North America. This graph shows the growth and decline of the Modular Hornet/OmniBER products from 1994 to 2002:

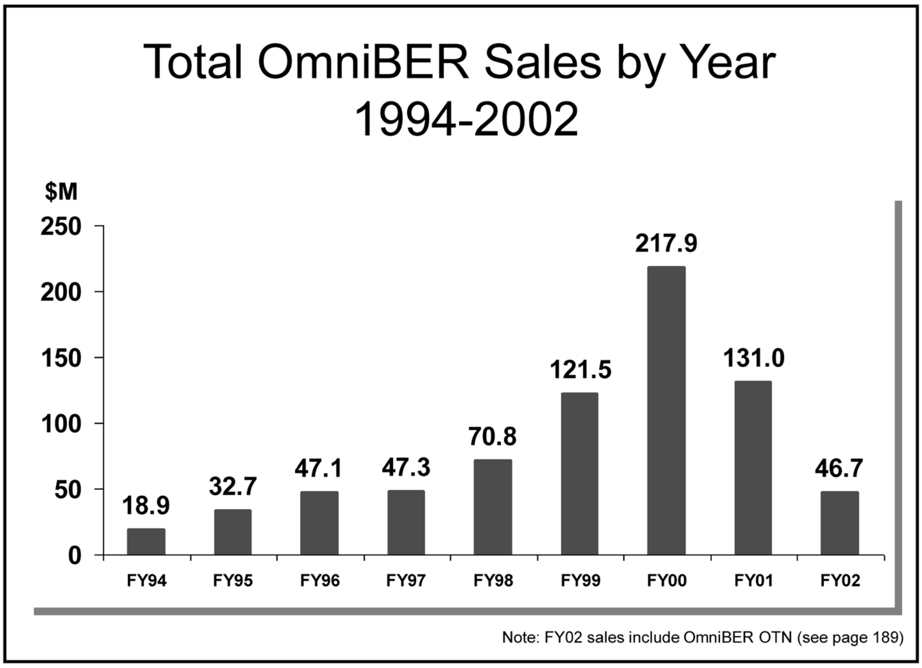

The end of the road came in 2002/2003 when the market died. The unit preserved in the National Collection is a lab prototype of the 37718A made around 1997/8. It differs very slightly from production models, but demonstrates the high quality of the prototype instruments built at that time.

Overall the Modular Hornet family of Digital Transmission Analyzers shipped over 14000 units and created total revenue of around $800M during a 10 year period, much of which occurred between 1997 and 2001. Probably 95% of this was exported, creating significant value for the Scottish economy. It was a fitting tribute to the tremendous skill and ingenuity of the engineers who designed and built these instruments at South Queensferry.

High Speed BER Testers

The chronology of this account now takes a step back in time to the late 1980s. In terms of application, it also takes a step back even further to the very first entry to the market in the early 1970s with the 3760/61, described in Chapter 9. This pioneering test set was aimed at the R&D activity for the first digital transmission systems over coaxial cable. Now, 15 years later the market had moved on to fibre optics.

In the late 1980s, fibre optic transmission was rapidly developing and 400 Mb/s transmission systems were commonplace in North America. In 1988 AT&T introduced the FT Series G system running at 1.7 Gb/s, while Fujitsu in Japan had a 800 Mb/s system. Like "Moore's Law" in computing, every year or two saw another huge leap forward in fibre optic bit rates and the spans possible with the new single-mode fibre operating at longer wavelength of 1550 nm. A great deal of research went into developing high-speed narrow-band laser transmitters and optical receivers and the associated high-speed electronics using Gallium Arsenide technology. This was moving rapidly into the Gigabit per second region, way beyond anything South Queensferry had so far produced.

The high-speed test equipment market was dominated by the Japanese, particularly Anritsu. In the 1980s, HP entered the optical measurement market, describing it as "Lightwave", linking it to its market-leading microwave instrumentation. This multi-divisional initiative was organised through the Photonics Measurement Strategy Council (PMSC) around 1986/87. Geoff Waters recalled this initiative which was chaired by Bob Allen, corporate engineering manager,

"*I spent a considerable amount of time over three years or more with a group of managers from Germany, HP Labs and the Microwave Divisions in Santa Rosa, who had the remit to move HP into the photonics business. Bob Allen never got much credit for this, but he should have.*"

Within HP, the pressure was on to produce a high-speed pattern generator and error detector to complement this and compete with the Japanese.

Earlier, South Queensferry had done some investigation on a 600 Mb/s BERT, presumably to test 565 Mb/s transmission systems (4 x 140 Mb/s in the old PDH system), but this was quickly overtaken by optical developments. In fact various "high-speed" investigations had been undertaken by the R&D team over the years starting with the 300 and 600 Mb/s project in the early 1970s, however no complete product materialised. The expertise was certainly there to design such a system, but maybe the market demand wasn't. There were some 565 Mb/s transmission systems around in the mid-1980s, however the economics weren't that good because of the high-frequency signal loss or attenuation on coaxial cables[6]. 140 Mb/s tended to be the optimum rate, as it was also for digital microwave radio. The economics changed radically when single-mode optical fibre with its enormous bandwidths entered the market in the mid-1980s. Almost overnight, transmission rates of 1000 Mb/s (1 Gb/s) and above became possible and economically attractive.

---

[6] A 565 Mb/s signal on coaxial cable using a 4B3T line code would have a top frequency of around 400 MHz requiring very close regenerator spacing. For a capacity of around 8000 telephone channels, this was a poor use of plant compared to the analogue 60 MHz FDM system handling 11,000 channels.

71603A/B 3 Gb/s Pattern Generator and Error Detector

Finlay Mackenzie, General Manager, recalls that he and the then R&D Manager, Robin Myles, visited Collins Radio in Dallas in the late 1980s to get their views on a 1.2 Gb/s BERT, however other customers in Europe were asking for 2.4 Gb/s. The Division settled on the higher bit rate. Thus work started on a new product that eventually became the 71603 3 Gb/s pattern generator and error detector.

Its feature set would be remarkably similar to that of the 3760/61 pattern generator and error detector introduced 15 years earlier. Like the earlier instrument, the new product targeted the parametric test of transmission components and electronics, and now optics. The main difference was the bit rate that would be 20 times higher, an indication how far digital communications had come during the intervening years. The new instrument had simple binary data and clock interfaces, with PRBS and digital word pattern generation. While the first BERT 15 years earlier had a 10 bit programmable word pattern, the 71603 could do 8000 bits and in the later B-model version, 4 Mbits! Like the first BERT, there was a facility for setting relative delay between clock and data, and the new error detector could automatically align clock and data timing for optimum results, a necessary feature for parametric test.

The 71603 was designed as modules to plug into the Modular Signal Analyzer (MSA) or Modular Measurement System (MMS) mainframe designed at HP's Santa Rosa Division in California. Geoff Waters recalled this was a result of the PMSC, "*The choice of MMS came directly from the Council, as several new photonics products were using this system. Santa Rosa also developed some optical/electrical and electrical/optical converter modules which were used with the new BER Testers.*" MMS was intended for high-performance microwave and lightwave instruments so was ideal for the new Gigabit BERT, although it was viewed by the South Queensferry design team as a rather expensive platform. The mainframe provided a colour CRT display with keypad and softkeys, power supply and cooling. The analyzer comprised three modules: a pattern generator (70841), an error detector (70842) and a 3.3GHz clock synthesizer (70311). The synthesizer was developed as a joint teaching project with Napier University in Edinburgh.

The high data rate meant the PRBS and word patterns needed to be created at a lower rate and multiplexed up. In this case, a 1/16 rate of up to 200 Mb/s was used so that the pattern generation could be implemented with bipolar gate arrays. The 16 parallel streams were then interleaved through a specially designed multiplexer to generate the full-rate signal. Similarly, the error detector used a demultiplexer to generate 16 parallel streams processed in a bipolar gate-array. Both multiplexer and demultiplexer were fabricated using a high-speed proprietary HP bipolar silicon process. The latter stages and the output amplifiers had to be very high bandwidth

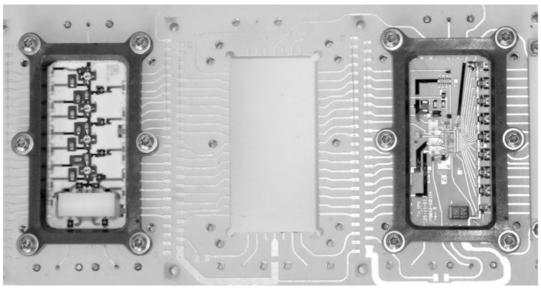

and were constructed in South Queensferry's Technology Centre. An example is shown here. Output waveform rise time was less than 90 picoseconds.

The 71603A Error Performance Analyzer was introduced in summer 1990 and had the in-house codename "Daedalus". It was superseded by the B model, codenamed "Theseus" in 1992. Over the 10 year production life, the system shipped around 1150 units at an average price of about $110k, producing $120M in revenue. A 1 Gb/s version (71601A) was marketed for a short time but made small sales and was discontinued.

71612A/B/C 12 Gb/s Error Performance Analyzer (Unit Preserved NMS T.2010.73.1-2)

Tracking the rising bit rate of optical transmission, South Queensferry developed the 71612 12 Gb/s pattern generator and error detector. It was introduced in 1993 and put HP into the market with one of the fastest digital test sets at the time. The initial target was the new generation of undersea fibre optic cables which were the first systems to operate at 10 Gb/s. Early sales were into companies such as AT&T, Alcatel and BT Labs, developing optical submarine repeaters. Later the market also included terrestrial systems with the introduction of 10 Gb/s SONET/SDH (OC-192/STM-64) in the late 1990s, when the market "boomed".

The system was again developed as part of the MSA/MMS platform like the 71603, although in this case, the pattern generator and error detector were built as two horizontal full-width trays which fitted into a standard instrument mainframe. This was interfaced to the MMS display for control and this display mainframe also housed a microwave synthesiser clock module (70340). The measurement mainframe had the product number 70843. With extreme pressure from key customers and competition from Anritsu, the development was completed in under two years, a significant achievement given the many design challenges and the new technology.

The feature set was pretty well identical to the earlier 3 Gb/s unit with highly flexible voltage settings on the binary and clock interfaces and variable delay. The programmable word pattern was further increased to 8 Mbits. The rise and fall times were specified at less than 30 picoseconds, and were typically around 20 ps, as this photograph of the waveform at 10 Gb/s shows.

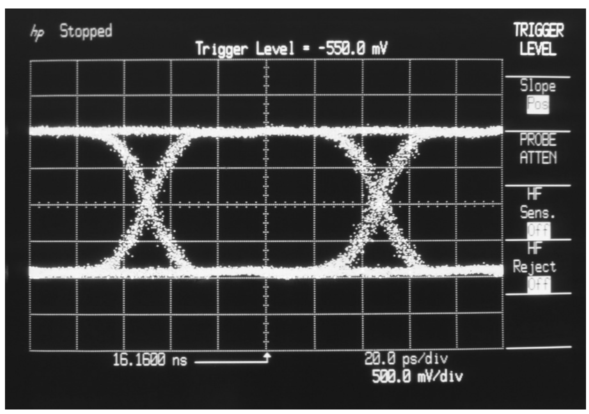

The design team used quite a cocktail of very high-speed gate arrays and other technology from around the world to build this advanced product. Geoff Waters recalled a global technology tour to source components, "*I visited HP Labs, HP microwave divisions, Vitesse in California, NEL in Japan and a host of other semiconductor manufacturers. At NEL I saw their state-of-the-art logic ICs, some of which I'm sure were already operating* *at 15 Gb/s. I arranged for samples to be obtained.*" The pattern generator used 1/32 rate multiplexing with the 32 parallel pattern generators, half in each of two ASIC[7] gate-arrays from Vitesse, running at up to 400 Mb/s. The PRBS patterns were generated internally to the ASIC, and eight random access memories per gate array provided the 8 Mbit programmable word function. The first stage of multiplexing (4:1) took the rate to 1.6 Gb/s using ECL devices and created eight sub-rate streams. These sub-rate streams were multiplexed up to full rate using modules manufactured in Japan by NEL (a subsidiary of NTT) using their market-leading Gallium Arsenide chips (source-coupled FET) connected by ultra-fine coaxial cables.

This final stage of multiplexing was very high speed and to preserve the integrity of the signal timing, the chips were positioned in a circular hole in the printed circuit board with a "starburst" format for connections. The final output amplifiers were designed at South Queensferry based on Gallium Arsenide travelling wave amplifiers developed by HP in the USA for microwave applications, with manufacture done at South Queensferry's Technology Centre.

The error detector used much of the same technology in reverse, ultimately making measurements at 400 Mb/s. A significant challenge was the provision of variable delay between clock and data. This was accomplished using modules produced by Wessex Electronics in Bristol to South Queensferry's designs. South Queensferry had worked with this company for a number of years as they had also supplied the variable delay/phase shifter module for the earlier 3 Gb/s BERT. A novel feature in the error detector was a function called "Error Location Analysis" which allowed the user to locate the exact position of the error in a pattern, helping determine the cause of pattern-dependent errors

The complete system had a list price of over $400k, though usually sold for $330k to $350k after discounts. It was introduced in October 1993, and had the in-house codename "Ariadne", continuing the Greek mythology theme of the Gigabit BERTS. In the early years, sales ran at 3 or 4 per month, but in the late 1990s the market went crazy. Orders for up to 30 units at a time came in from equipment manufacturers like Nortel. Geoff Waters, responsible for sales development, remembered having a weekly call to a Nortel director to schedule deliveries. Nortel even asked if HP could set up a separate production facility for them!

Geoff Waters remembered this bizarre period:

"*In February 2000, I already had a spreadsheet compiled by Nortel for their requirements for the year. The forecast was 44 units, but we actually got orders for 86 units which showed how much Nortel had increased their estimates for 10 Gb/s test stations, which used the 71612.*"

---

[7] Application Specific Integrated Circuit FET – Field Effect Transistor

It was all part of the frenzy of the "dotcom bubble" when masses of optical fibre was laid everywhere, much of it "dark fibre" that wasn't brought into service until years later. In 1999, orders rose to 165 units. In 2000, the Division couldn't make the instruments fast enough when orders reached 413 units, and revenue was $120M. It was as if five to ten year's orders had been received in one year! These were exciting times.
John McElroy recalled that a cross-functional team was organised to raise manufacturing output tenfold, from two units per month to 20. Bob Thomson, who had been a design engineer and project leader in R&D, was part of the team:

"*My last four years were spent in manufacturing on High-speed BER, and it really was an excellent experience, some of the best time in my career.*"

The Division clearly thought this bonanza would continue – the forecast was 553 units in 2001 and 719 in 2002! It wasn't to happen. Instead, as fast as it had mushroomed, the demand died away again as the bubble burst. Geoff Waters remembered that the business in 2001 was very erratic with still occasional large orders but also many cancellations. In 2002, months went by with no orders. The market was flooded and further sales became very difficult. The 71612 was discontinued in 2003. Returning to Greek mythology, Ariadne was the granddaughter of the sun god Helios, and for a short time South Queensferry's "Ariadne" blazed like a sun – a dying sun. It was the test set of choice for the makers of optical transmission that fuelled the "dotcom bubble".

During its 10 year life, the product underwent some minor upgrades and some new applications and measurement methods were developed, including raising the maximum bit rate to 12.5 Gb/s. Some PC application software was introduced in 1996 to work with the 71612A. The E4543A provided automated analysis of the digital eye-opening by using the programmable data versus clock timing facility on the error detector. Another package, the E4544A, simplified the programming of the 8 megabit word capability for simulation of the 10 Gb/s SONET/SDH frame and other applications such as 10 Gigabit Ethernet. Both packages sold for $10k. The upgrades to the 71612 were designated as the A, B and C models. The B model was introduced in January 1999 and was the big seller. In total some 700 systems were shipped worth $250M in revenue.

The system preserved in the National Collection is a 71612B (70843B) with serial number 487. It was most likely manufactured in 2000/2001. Associated with the full working unit, additional pattern generator and error detector modules have been preserved at the Museum to display the technology used.

The End of the Line

The Modular Hornet instruments, OmniBER 717 and 718, produced with such success in the 1990s, targeted both the development/production market as well as network operators installing and maintaining transmission networks. By the late 1990s, several competitors introduced more compact tablet-like instruments or products that looked a bit like laptop PCs. One company in particular, Digital Lightwave based in Florida, introduced a product called the Network Information Computer testing PDH, SONET and Ethernet networks. Using the latest advances in component miniaturisation and greater integration, these new products were considerably more compact than the Modular Hornet, and were seen as desirable for field use by network operators.

HP had become interested in this network operator market too, and several divisions started to introduce products aimed at this business. At South Queensferry, the product development and marketing teams were split into two groups, one focussed on the R&D and production market (RDP section) and one focussed on the digital network performance (DNP) business. In 1999, HP's Test and Measurement business was spun out as Agilent Technologies, HP from then on being exclusively an IT and software company. So, these final products were branded Agilent, although designed and built by the same people in the same place.

E7580A Prober 2 Handheld E1 tester (Unit Preserved NMS T.2010.75)

The first new entry into the DNP sector was the handheld ProBER 2, introduced in 1999. It did much the same job as the successful 37722A EI Test Set in the Hornet package, which it replaced. Also, the Division had marketed for several years a handheld 2 Mb/s tester (37742A) which was a badge-engineered product from ICT in Spain. As mentioned earlier, these products had done well which led to this new development. The ProBER 2 targeted the digital services and private network market of the time, which was predominantly at 2 Mb/s (E1) and below.

Like its predecessors, it could fully test the multiplex frame structure. The ProBER is interesting as it was the first complete product South Queensferry developed under contract on-site but not in its own R&D lab. The development was done by a company, ECS Telecom, run by ex-HP engineers and based in rented space at the South Queensferry site, working closely with Peter Scott from R&D. They had previously done contract work on some modules in the OmniBER 717/718. The product came in a smart carrying case with battery charger and cables.

It cost $4.7k and it is thought 3000 to 4000 units were shipped until 2003, when it was discontinued. This may have been part of a decision by Agilent to get out of this market sector, although the 2 Mb/s market was by then gradually being overtaken by Ethernet and Broadband Access. There were plans for follow-on ProBER handhelds but they didn't go ahead which was probably a mistake as competitors such as JDSU did good business in this area in the years that followed. The unit preserved in the National Collection has serial number 125 so may be quite an early unit.

J2126/27A Transmission Test Set (Unit in Preservation)

In a way, South Queensferry's flagship product for the new approach to the network operator market was the J2126A and J2127A family of transmission test sets. Taking advantage of the latest technology in gate-array ICs and SMT, the new products were very compact considering the measurement capability. A measure of how far this had gone was that the functions of a Hornet product from 1992, the 37724A SDH/PDH Test Set, were now all contained on one relatively small board in the new instruments. Part of this board is shown on the following page to illustrate the miniaturised construction (see page 170 in this chapter for comparison):

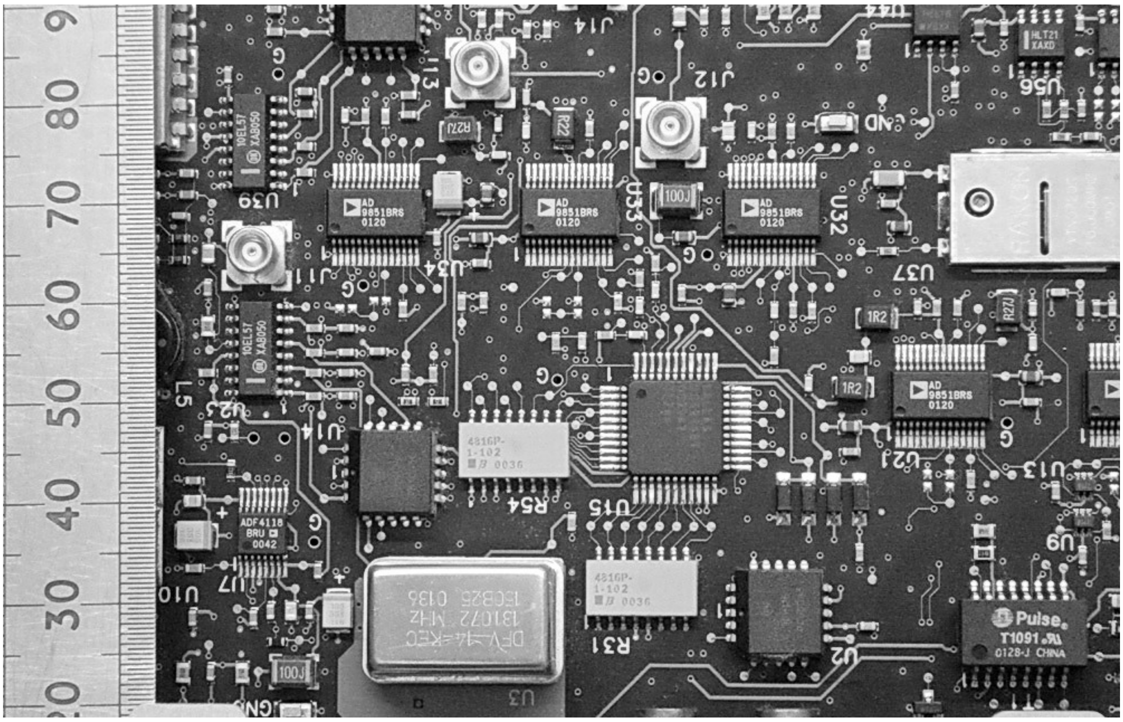

Development work started in 1999, and completed prototypes were ready in 2001. A former marketing engineer on the project, Milton Gilmour, recalled later the remarkable speed of this development:

"*From memory, the J2126A project was planned as requiring 18 months to complete. It was actually delivered two-months early. For a product of its complexity, that was designed from the ground-up (hardware platform, all functional electronics, and internal platform software), this was an awesome achievement. To me, it illustrated the truly world-class engineering skills that existed in the SQF organization at the time.*"

The J2126A, codenamed "*Lynx*", had a striking appearance. A local firm of industrial design consultants helped the in-house mechanical designers develop an attractive and compact package with a nicely integrated front-panel layout. Around 60% of the front panel was occupied with a 9" diagonal back-lit colour LCD panel, while the keypad borrowed some ideas from the OmniBER 718 product.

Mechanically the unit comprised three sections. The front panel contained all the processor and memory hardware, and this module plugged into the main motherboard at the base of the instrument. Into this plugged the measurement hardware with the optical and electrical measurement connectors on top of the instrument, similar in concept to the Modular Hornet platform. The power supply and cooling module at the back of the instrument plugged in to the back of the motherboard. Thus, by removing a few screws and a couple of tie bars, the instrument could be disassembled into the three sections. It was one of the best bits of mechanical design done at the Division.

The J2126A “*Lynx*” was the smaller configuration, providing test up to 2.5 Gb/s SONET and SDH as well as PDH. A cheaper version going to 622 Mb/s was available by deleting the 2.5 Gb/s measurements. The J2127A was similar except it had a longer mother board and thus a deeper instrument to accommodate another module which would typically be a 10 Gb/s optical interface. Alternatively, the extra slot could be used for a Gigabit Ethernet module introduced in 2003. This extended version had the codename “*Cobra*”. A further extension to the J2127A, codenamed “*Tiger*”, was planned but never taken forward.

The objective was to get cost down and also improve ease of use. With a software feature called “*Signal Wizard*”, the instrument could automatically identify the properties and frame structure of the incoming signal. This and other features were marketed as “*Extreme Productivity Improvement*” or XPI to use the obligatory three letter acronym.

Launched in 2001, the product had an auspicious start with advance orders of $30M, many for the 10 Gb/s version. The marketing launch of integrated literature and packaging looked very professional and won the Agilent “Best Marketing Award” in 2001.

Sadly, this promising start didn’t last as the global market began to collapse. This innovative product entered the market at the worst possible time and over the next three years, sales never ran above a few hundred units a year. In 2004 sales amounted to 113 units and the product was discontinued. Overall sales were 1065 units with total revenue of $56M. Product price varied from $120k for the 10 Gb/s J2127A, to around $25k for the 622 Mb/s J2126A. A J2127A “*Cobra*” is preserved.

J7230A OmniBER OTN 10G (Unit Preserved NMS T.2010.75)

This was the final development of BER testing at South Queensferry, and ran more or less in parallel with the “*Lynx*” project, sharing technology. Whereas Lynx was aimed at network operators, the J7230 targeted the R&D and production market. It filled gaps in specification of the J2127A and added jitter measurements. The title OmniBER OTN made reference to the new ITU-T standard for optical transmission, Optical Transport Networks (OTN).

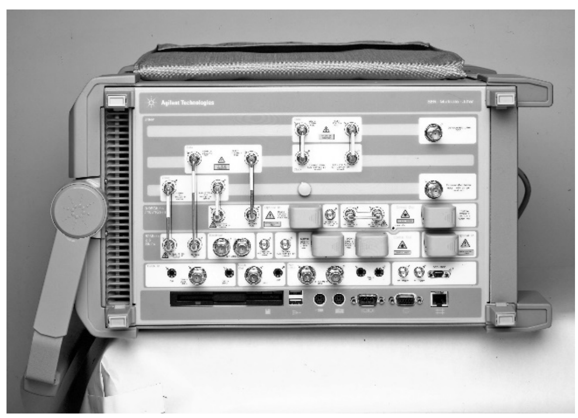

Built into a larger box than the J2127A, the OmniBER OTN was configured with the measurement modules as horizontal blades plugging-in from the side with the signal connections on the right side panel similar to Modular Hornet. The main additional feature of the OmniBER OTN was the optional jitter capability to 10 Gb/s which occupied the extra space in the instrument. Three main versions were sold: the J7230A/B which did BER up to 10 Gb/s (codenamed "Panther"), the J7231A/B which added jitter to 10 Gb/s (codenamed "Merlin"), and the J7232A which had a top rate of 2.5 Gb/s BER and jitter. The price of the 10 Gb/s instrument was around $150k.

Introduced in 2002, after the market collapsed, the sales of OmniBER OTN were dismal, despite some Herculean attempts at sales development by the marketing team. Between 2002 and 2005, all the variants only managed total sales of 440 units with revenue of $55M. It was all downhill and by early 2005 things were so dire the decision was taken to discontinue the products.

Ian Johnston recalled that the last hardware project at the Division was the development of an add/drop Ethernet module, codenamed "*Theoden*", for the J7231B. "*As far as I remember, the printed circuit board assemblies were out getting fabricated when the 'coup de grace' was delivered!*"

It was the end of the Telecom Division too, which was by then making heavy losses. The operation was closed down in 2006, almost exactly 40 years after the whole enterprise was set up. Production activity ceased on 26th January 2006, when the last 71612C was shipped following factory repair.

The J7230 OmniBER OTN unit preserved in the National Collection thus represents the final evolution of South Queensferry's digital transmission test sets, a history stretching back nearly 35 years. Years later, a competitor asked a former marketing engineer why Agilent had obsoleted the best SONET/SDH tester in the world. The answer was simple: the Company couldn't afford to make it any more.

During the last year or two a couple of further developments took place. A 40 Gigabit version of OmniBER OTN (codenamed "*Quantum*") was created, anticipating the move to higher rates in the network. Designated the J7230C it was planned to sell for around $350k, however it never got beyond the prototype stage. Two prototypes were eventually donated to Glasgow University. However, there was a spin-out from this through

Queensferry's long association with the HP telecom operation in Melbourne, Australia, starting 20 years earlier with 75000 Series 90, as described previously.

The Australian Telecom Operation had a particular interest in broadband, initially in Asynchronous Transfer Mode (ATM) cell transmission, and later in Internet Protocol (IP). They had success in the late 1990s onwards with a product called the RouterTester, which combined optical transmission with IP packet generation and analysis. The 40 Gb/s optics assembly from the "*Quantum*" project was adapted and enhanced to form part of this product range, which was marketed as the N2X line, Melbourne adding a large packet processor behind it.

The 40 Gb/s (OC-768c)[8] packet processing module for testing IP routers became the E7320B, and Ian Johnston recalls building E7320B optic modules at South Queensferry up until 2009. Some work also continued at South Queensferry on software enhancements under a project called "*Schiehallion*". Without "*Quantum*", this flagship product would not have happened, or at least not as quickly as it did.

Another development in the early 2000s was a multi-port version of OmniBER intended for production test. Called OmniBER XM, J7241-5A, it was unfortunately afflicted by the same market malaise as the other products. However, as Andrew Wilson from product marketing pointed out, the definition of the products may have limited their success too.

"*They were intended really for just one narrow application – testing path-level protection in multi-service platforms. The flaw was that this application was too specific, so the XM never got the deeper virtual-concatenation test functions of the OmniBER OTN that would have made it successful.*"

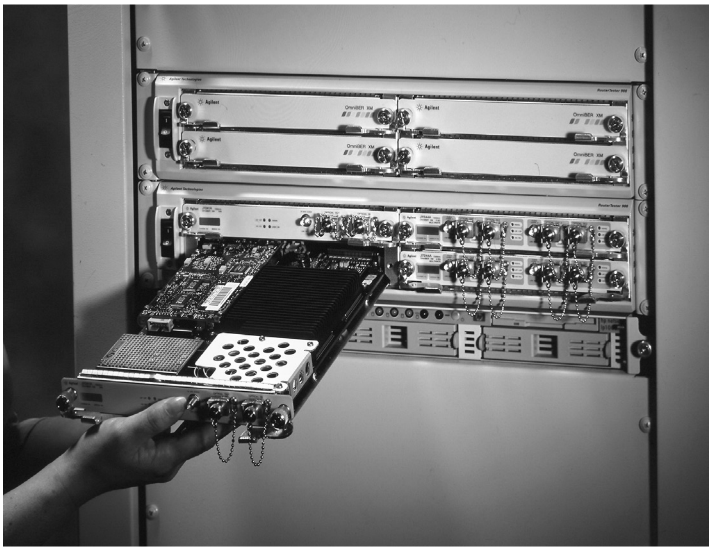

There were effectively two products, the J7241/2A operating at 10 Gb/s codenamed "*Yellowstone*" and the J7244/5A for 2.5 Gb/s, codenamed "*Nugget*". The OmniBER XM range also became part of the N2X line mentioned above, and product responsibility eventually transferred to Melbourne.

In 2009, around three of the remaining staff at the factory, and a few from Melbourne and a related operation in Vancouver, Canada, transferred when the Agilent Corporation sold the whole business to a long-time competitor, Ixia of Calabasas California, for $44M cash. It was the end of the road for Digital Transmission Test at South Queensferry.

[8] OC-768 was the North American nomenclature for the SONET rates – for example OC-3 was 155 Mb/s, OC-48 2.5 Gb/s, OC-192 10 Gb/s and so on. The "c" in OC-768c indicates that the payload area in the SONET frame is not sub-divided but is "concatenated" to provide one wideband channel for conveying IP packets. Gigabit Routers from companies like Cisco used concatenated interfaces.

Conclusions

The last few years were indeed a depressing end to a long and successful history of developing and producing digital transmission test sets at South Queensferry. It was in stark contrast to the booming success a few years earlier. The business just faded away and the product line and Division died with a whimper. All the accumulated technical expertise and market knowledge gained over the years drained away. It is something of a mystery why Agilent couldn't find a use for all this brilliant engineering capability, and just let it go. It seemed like the Company had lost interest.

In the end, South Queensferry's Digital Transmission Analyzers, and the Division itself, were caught in a trap from two directions. There was the dramatic fall in sales following the collapse of the dotcom bubble, which undermined the Division's finances. But also the market had started to change. The emphasis in the years following 2001 was less about transmission parameters, error rates and transport overhead, and more about the actual payload area of the SONET/SDH frame. Here the focus was on packet processing, traffic loading and statistics in the new Ethernet and Internet age. This was the business of the Australian Telecom Operation and its Router Tester, rather than South Queensferry.

On a positive note, much of the 35 year history was a success story, and the BERT family product line created secure and well-paid jobs for hundreds of employees, and some spent their whole career with this business. The product range included some superb engineering achievements. Overall it generated more than $2B in revenue, and with exports of over 90%, created a great deal of value for the Scottish economy.

Acknowledgements

Thanks are due to the following South Queensferry employees who have contributed information and recollections to this chapter and the previous one:

Finlay Mackenzie, Tom Crawford, Robin Myles, Ivan Young, Bob Thomson, Peter Scott, Jim Kendal, Geoff Waters, Jim Barron, John Coster, Stuart Connelly, Martin Curran-Gray, John McElroy, John Wotherspoon, Milton Gilmour, Ian Johnston, Andrew Wilson and David Robertson.

11

Chapter Eleven

Signalling Test Sets

Signalling in telephone systems is the method by which calls are set up and terminated between terminals through a network of switches and transmission links. Telephone systems up until the 1960s relied on electro-mechanical switching using relays, step-by-step "uniselector" rotary switches and two-motion (vertical and rotary) selectors, known as the Strowger exchange. The telephone typically had a rotary dial that sent out a specific rate of pulses to the exchange by interrupting the line current supplied from the exchange. This matched the reliable switching rate for the selectors.

A call was started by lifting the telephone receiver which alerted the exchange (by sending a DC current) and a "line-finder" uniselector clocked round to determine which line was "off hook". Once connected, the exchange sent back "dial tone" to the calling telephone, and the user could then start dialling the number using the rotary dial. In the simplest implementation, these dial pulses operated the first stage of the switching, and the next set of pulses passed through the connection established by the first switch to control the next stage, and so on. There was no central control, the dial pulses set up the path progressively through the network, stage by stage, and this electrical connection also carried the speech signal and sent back to the calling party signals such as "line busy", "number unobtainable" and the "ringing signal" from the far end.

In effect, the control function was distributed through the switching network. Later, more sophisticated Strowger systems, for example those used in early Subscriber Trunk Dialling (STD), first stored the dial pulses in an electro-mechanical register, which gave the advantage of more efficient use of the actual switching equipment. However, with all these systems the speed of setting up a connection was governed by the pulse rate from the telephone dial, nominally ten pulses per second.

Large Strowger Exchanges (or Central Offices as they are called in North America) were impressive machines particularly when handling a high traffic load, and despite their ungainly size and old-fashioned technology they were probably the first large-scale application of logic theory[1] and traffic statistics. The statistics came about because it is uneconomic to make a telephone switch completely "non-blocking", so a trade-off is made between level of service and cost. Even with modern telecommunications equipment such as digital telephone switches or Internet routers, if the traffic is very heavy, some customers may be denied service or "blocked". Although the theory was sophisticated, the system "programme" was completely hard-wired into the switches so it was impractical to make changes or alter services.

Fast-forward now to the 21st Century, when we expect the Public Switched Telephone Network (PSTN) to offer caller ID, call forwarding, free-phone (800 toll-free), premium rate calls, number translation, call barring, voice-messaging and so on. This is down to the fact that modern telephone switches are all controlled by computers, once referred to in the industry as Stored Programme Control (SPC). In contrast to the old Strowger exchange, the control is now completely separate from the actual path switching; in fact it doesn't even need to be co-located.

Perhaps even more dramatic is the growth in mobile telephones, with billions in use around the world. Very few users have any idea of the complexities of keeping track of where mobiles are, so that they can be called anywhere in the country or even the world! This is only possible because of a huge network of hidden computers, exchanging information on location, billing and available services, so that the connections can be set up between users in both fixed and mobile networks.

This extraordinary evolution started in the 1960s with the introduction of a new generation of analogue telephone switches called "crossbar". In these switches, the old Strowger step-by-step rotary switches were replaced by a matrix of reed-relay switches, with a reed relay interconnecting an input and output line at each intersection in the matrix – for example 144 reed switches for a 12-input by 12-output matrix. The reed switch is a delicate cantilever contact in a glass tube which is inserted in an electromagnet. The reed switch can operate in a few milliseconds and perhaps as little as a millisecond. Using this technology, a crossbar switch could set up a connection in the same time that a Strowger exchange could accept just one pulse of dialling. Because the switching element was hermetically sealed, it was also much more reliable than the older switches. The reed relays were activated by control logic from a Master Control Unit (MCU), so for the first time the control of the switch was centralised and separated from the switching fabric. Although the control was digital, the switch itself was still an analogue metallic path and so was compatible with the earlier systems.

The first crossbar exchange was introduced in the USA in 1965 (the Bell System 1ESS standing for Electronic Switching System), while in the UK the first crossbar exchange (TXE1 standing for Telephone eXchange Electronic) was brought into service in 1968. Over the next decade, the system was developed with more powerful Stored Programme Control using microprocessors. The final UK development was the TXE4 (handling up to

[1] It is interesting that the first electronic computer, the Colossus used for code-breaking at Bletchley Park during WW2, was designed by Tommy Flowers who had worked on telephone switching at the Post Office Research Labs at Dollis Hill in the 1930s. After the War, Flowers returned to Dollis Hill and worked on the first digital computer-controlled exchange which was installed on an experimental basis in Highgate in London in the early 1960s. It was not particularly successful, mainly because the electronic technology of the day wasn't up to the job, however it pointed to the future direction.

40,000 lines) of which several hundred were installed until the 1980s, when crossbar was superseded by the System X digital exchange[2], based on Time Division Multiplex (TDM) and PCM (Pulse Code Modulation) encoded voice. Although the switch fabric was now digital and electronic, the concept of the Stored Programme Control (SPC) remained the same.

The remainder of this chapter focuses on the signalling system between these SPC computers, the signalling network and the messages that are used. South Queensferry's signalling test sets were able to decode and analyze the multitude of complex messages and also generate signalling traffic for systems test.

Channel Associated and Common Channel Signalling

The major division in signalling systems is between Channel Associated Signalling (CAS) and Common Channel Signalling (CCS). Today, common channel systems predominate outside the local loop, but the technique is by no means new. In the days of manual exchanges, "Order Wire Working" was a form of common channel signalling, whereby the operators spoke over one circuit to set-up calls on a larger group, thus eliminating the need for expensive signalling equipment associated with each voice circuit.

Subscriber Trunk Dialling (STD) was introduced in the UK in 1958[3] and was widespread by the late 1960s between all major towns and cities. The equivalent in North America was called Direct Distance Dialling. This meant a subscriber could dial long-distance directly, rather than needing to use an operator when dialling outside their exchange area. Long-distance connections were over multiplexed transmission systems such as FDM coaxial cable and later digital transmission, so it was necessary to send the dialling information to the distant exchange over the transmission facility.

To do this, tones were sent over the telephone channel as the call set-up progressed, similar to the Dual-Tone Multi-Frequency (DTMF) dialling used on the subscriber access line today. These tones were decoded at the far end and translated into dialling pulses for the far-end electromechanical exchange. This process was called Channel Associated Signalling (CAS) as the signalling information was carried by the telephone channel or was closely associated with it as the call was set-up. The standard used in the UK and Europe was CCITT[4] R2 (Regional 2). A development of this standard was CCITT Signalling System No.5 which was used both nationally and internationally. Like R2, it used tones which were sent directly over an analogue telephone channel or encoded in PCM on a digital facility. Various forms of digital CAS also evolved in the first generation of digital systems. One used in North America was called "Bit Stealing" where the least significant bit of certain 8-bit PCM words was intermittently "stolen" to create a low-speed CAS data channel carrying signalling for that particular telephone channel. The missing bits were undetected on a voice call, but it did mean that when the channel was

---

[2] Similar digital exchanges based on Time-Space-Time switching were introduced by major manufacturers such as Nortel (DMS), Alcatel, Ericsson (AXE), Siemens (EWSD) and AT&T (Lucent) with their well-known 4ESS and 5ESS systems, capable of handling over 100,000 lines.
[3] The Queen made the first official STD call from Bristol to the Lord Provost of Edinburgh on the 5th December 1958. See http://news.bbc.co.uk/1/hi/uk/7766631.stm
[4] International Consultative Committee for Telephone and Telegraph, now ITU-T, International Telecommunications Union – Telecommunication standards.

used for data, only seven bits could be used, hence the 56 kb/s (rather than 64 kb/s) rate common in the USA. In the European E1 (2 Mb/s) primary multiplex frame carrying 30 PCM voice channels, one timeslot (TS16) was allocated to carry the signalling data for the associated 30 telephone channels.

As SPC exchanges proliferated in the 1970s (initially crossbar and later digital), signalling between exchanges increasingly moved to dedicated data links between the SPC computers, and signalling was no longer directly associated with the physical telephone circuits. This became known as Common Channel Signalling (CCS), and the first large scale application was CCITT Signalling System No.6 (SS6). It was mainly deployed in North America from the mid-1970s onwards, where it was known as Common Channel Interoffice Signalling (CCIS). Here, "common" denotes the collective handling of signalling for many voice channels.

SS6 had been developed in the late 1960s and early 1970s, before the advent of the Open Systems Interconnection (OSI) layered model for messaging protocols, which didn't emerge until the late 1970s. SS6 was therefore an early data-communications protocol designed in a monolithic way for a specific purpose, and consequently was more difficult to modify and adapt to new applications. This limitation was solved by the next CCS evolution, Signalling System No.7.

Signalling System No.7

Initial development took place in the 1970s, and the first CCITT standard appeared in the 1980 Yellow Books, with further refinements and additions in the 1984 and 1988 CCITT Red and Blue Book publications. Signalling System No.7 is known by a variety of abbreviations including SS7, CCS7, CCITT No.7 and C7.

Part of the definition of SS7 was the independent network which was developed to carry the signalling messages. This involved some key elements as shown in this simplified network diagram:

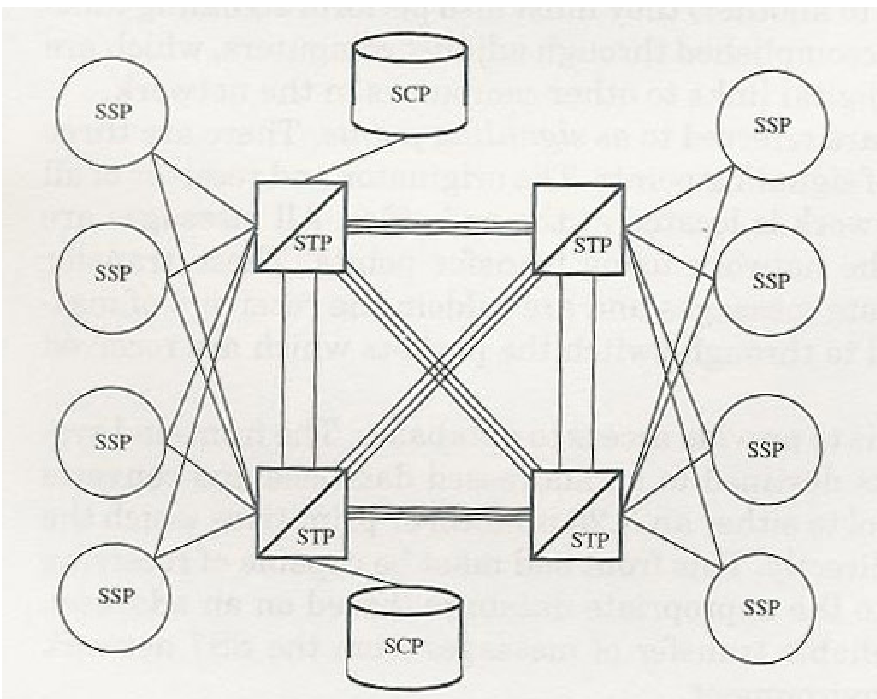

The Service Switching Points (SSPs) are the telephone exchanges running under Stored Programme Control as described earlier. The Signalling Transfer Points (STPs) are the data switches that route the signalling messages between the exchanges, and the Signalling Control Points (SCPs) are the databases used for telephone number translation (e.g. toll-free 800 numbers), billing etc. These are all connected by "Signalling Links", and the core network has multiple bi-directional links between STPs (as shown) called "Link Sets". Moreover, the STPs are usually arranged in pairs, cross-linked as shown. This high degree of network redundancy allows traffic load to be shared and signalling messages to be re-routed automatically if part of the network goes down. The SS7 network becomes the "nervous system" of the telecommunications network, so any faults could paralyse operations on a very large scale.

As mentioned, a major advantage of SS7 was that it uses a layered model or protocol stack closely similar to the OSI (Open Systems Interconnection) model defined slightly later. With a total of seven possible levels, each layer has a specific function that relies on the layers below and provides a service to the layers above. By carefully defining this functionality, SS7 could avoid the inflexibility of the earlier SS6, since any layer in the stack could be more readily modified or enhanced without incurring a lot of change in other protocol layers. This proved valuable as more advanced communication services developed, in particular mobile communications which was virtually unknown when the first SS7 standard was issued.

SS7 uses a close approximation to the OSI protocol stack, and is best considered in two halves:

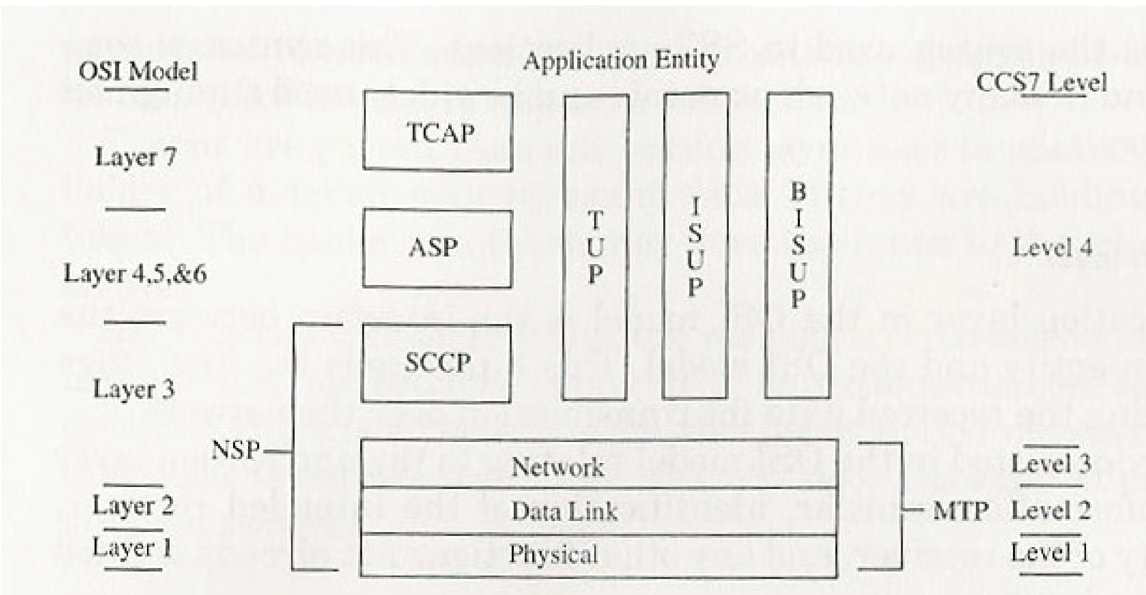

The lower half is called the Message Transfer Part (MTP), which corresponds to the first three OSI levels. Its function is to deliver the SS7 messages, reliably and without error to the correct destination. It operates in a similar way to the CCITT X.25 packet switched network introduced in the 1970s, and the Signalling Transfer Points (STPs) are packet switches. The packet is rather like an envelope which contains the message. The packet has the forwarding address (called a Point Code) of the node in the SS7 network (SSP, STP or SPC) known as the Signalling Point. It also has a sequence number so that the various packets can be placed in the right order at the receiving end as they may be out of sequence having been transmitted over different links in the link set described above. The MTP also does an error check[5] on the contents of the packet. If an error is detected, the packet is discarded and the transmit end is asked to resend it.

Inside the "envelope", the SS7 higher-level messages are encapsulated. This is called the Application Part (AP) or User Part (UP) and includes anything from a simple telephone call[6] to more complex datacommunications services. As SS7 evolved in the 1980s and 1990s, new applications were added, in particular various application parts for mobile-phone networks. Other applications such as the Transaction Capabilities Application Part (TCAP) are used for accessing the SCP databases.

In SS7, the packets are referred to as Signalling Units. SS7 uses three types. The most common is the Message Signalling Unit (MSU) which conveys the information between the exchanges and databases via the STPs in the packet-switched network. The second type is the Fill-In Signalling Unit (FISU) which keeps the system alive when there are no messages to be sent. The third type is the Link Status Signalling Unit (LSSU) which contains information on the status of the signalling links between nodes. The LSSU would initiate a rerouting of traffic if the quality on a link or link-set becomes inadequate.

This very brief overview of the SS7 network and messaging is only intended as a context for describing South Queensferry's signalling test sets. The topic is a large one, covered in detail in the CCITT/ITU-T Q.700 Series recommendations, although these are often modified in national and regional variants. There is also a very useful survey in the book "*Signaling System #7*" by Travis Russell[7].

Common Channel Signalling enabled many new services and networks, which would otherwise have been quite impractical to implement. At a more basic level, it meant that transmission and switching equipment could be used more efficiently[8] since traffic flows could be balanced more readily using the independent control path of the signalling network.

By the time the 1988 CCITT Blue Book Recommendations were issued, the SS7 standard was fairly stable, and switch manufacturers and operators adopted it for the signalling network on a global scale. Even the first generation SS6 in North America was gradually replaced by SS7. By 1999, the 1900 pages of standards in the 1988 Blue Books had expanded to nearly 9000 pages, a measure of the importance and the market for SS7 products and applications.

---

[5] This is a Cyclic Redundancy Checksum which is calculated on the packet contents on transmission and the recalculated for comparison at the receiver. If the answers are different, there is an error in the packet and retransmission is requested.

[6] Telephone call set-up initially used the Telephone User Part (TUP), though most operators adopted the more versatile ISDN User part (ISUP) for telephone operations.

[7] "*Signaling System #7*" by Travis Russell, McGraw-Hill 1995, ISBN 0-07-054991-5. Travis Russell was at the time a project manager for SS7 test products at Tekelec Inc., a major competitor for South Queensferry's signalling test sets.

[8] A good example is the classic case of the called subscriber being engaged on another call. In the old system, the telephone trunks would be seized stage-by-stage until the far-end exchange tried to connect and found the telephone was in use. It then sent an engaged tone all the way back to the calling subscriber and the call was terminated. In CCS, no trunks are seized until the signalling system has determined that all equipment en-route is available including the far-end telephone. If there is a problem, the local exchange simply sends back the engaged signal to the calling subscriber

South Queensferry enters the Signalling Test Market

Although signalling test had been of interest to the Division for some time, the market didn't look very promising, as Finlay Mackenzie (former General Manager) recalled.

"*With all the various types of signalling in use and the regional variations (17 in the UK), it was difficult to see how we could make a viable product. However, that all changed with SS7.*"

In the mid-1980s, the product development team sensed there was a rich seam for future development and the universality of SS7 meant an attractive global market. One individual, David Dack, described by some as a South Queensferry "visionary", undoubtedly championed this new opportunity as he had done with some previous innovations including the first use in HP of the Intel microprocessor in the early 1970s, the invention of the HP-IB Extender product, the development of a digital filter chip and the pioneering work on computer-based monitoring systems. As it turned out, the SS7 opportunity was far greater than anyone initially thought, even probably David Dack!

At the time, HP already had a high-end protocol analyzer, the 4955A from Colorado Telecom Division, that offered decodes of SS7 messages as well as other protocols such as X.25. So why did South Queensferry think the market needed another SS7 test set?

Although they do a similar job, there are two key differences between a general-purpose protocol analyzer and a signalling test set.

- The signalling tester needs to operate simultaneously with several parallel input streams because, as explained earlier, SS7 networks use a "link-set" between nodes consisting of several physical links sharing the signalling traffic. The tester needs to monitor several bi-directional paths simultaneously to capture all the signalling traffic.
- It also needs to do the message decoding in real-time so that it can perform functions such as call tracing and traffic analysis. For in depth analysis, protocol analyzers often store a data record and then analyse it post-capture, however this could easily miss important data in the constantly changing stream of SS7 messages.

The new South Queensferry test set was going to need some special-purpose hardware and firmware as well as analysis software. An investigation had already started when the Queensferry engineers came across a German company called Telenorma (a subsidiary of Bosch AG) who had developed some SS7 hardware and software.

Telenorma had been involved in design and production of SS7 equipment for the Siemens EWSD digital switch and had also developed a test system for monitoring and generating SS7 messages, which they needed for their SS7 work. Not being in the test equipment business, Telenorma approached the HP sales organisation in Germany with a view to selling the design, and this was passed on to South Queensferry. David Dack (R&D) and Finlay Mackenzie visited Telenorma and agreed the proposal looked interesting. Following a technical review of the hardware by David Guest and the software by Peter Locke, they assessed the system was functional but required some work to make it into a sellable piece of HP test gear. In July 1987, Finlay and Robin Myles (R&D Manager) went to Frankfurt to negotiate the contract. Telenorma would supply a working test system consisting of the signalling data receiver and transmitter controlled by an HP 9000 series 300 controller, source code for the software and firmware, and a package of consulting and support for a 12 month period. HP would have exclusive use of the design and the total cost of the deal would be 1.2M Deutsch Marks.

The deal was approved by Bosch, but as Finlay Mackenzie recalled:

"*Telenorma didn't understand why we couldn't sign the contract on the day. I called Dick Anderson and Bill Terry at HP Corporate in Palo Alto and they told me we would have to put the proposal to the Executive Committee at their next meeting. Finally we got the go-ahead, but we also had to involve David Baldwin (General Manager of HP UK). It all seemed to take ages and we began to worry that our German rivals Wandel & Goltermann would sniff out the deal.*"

In August, the South Queensferry team accompanied by David Baldwin and the HP lawyer, Robert Squibbs, returned to Telenorma and the contract was signed by Robin Myles on 20th August 1987. And so began South Queensferry's venture into the signalling market.

It has been suggested that the Division's interest in signalling had been triggered two or three years earlier as some basic channel associated signalling test had been incorporated in the large RATES system developed at South Queensferry in the early 1980s (see Chapter 7). The main customer for this system was British Telecom, and through them the Division was probably introduced to Samuel Welch, a world authority on signalling systems[9] and formerly head of signalling at the British Post Office Telephones. It is said he was the driving force behind the development of SS7. In the late 1980s, Mr. Welch (who was known to us simply as "Sammy"), visited the factory on a number of occasions and delivered lectures on signalling systems and telecommunications in general, as well as providing consultation for the signalling R&D team.

More than 30 years older than most of us and long retired, Sammy relished his role as the fount of knowledge. He would make his pronouncements, sometimes with pipe securely located in a gap provided in his lower rank of front teeth, waving his index finger back and forth: "*Error detec-shun-n-n by redundant-t-t coding; error correc-shun-n-n by re-trans-mishun-n-n. And this means this!*" Thus was his description of error control in SS7. Apparently he didn't hold the USA network in high regard:

"*A rate of Two Megabits-sh per Shecond-d, which is used by all the countries in the world, apart from one backward country I won't bother to mention.*"

Furthermore, the said Americans were "stupid" to have adopted the clearly inferior SS6 signalling. He was highly impressed with the lunches served in the South Queensferry cafeteria, and I remember on one occasion him commenting that at Rosyth Dockyard, where he did some work, they served "*Very good soup.*"

Armed with all this know-how and the basic tester from Telenorma, the R&D team was ready to start product development.

---

[9] His book "*Signalling in Telecommunications Networks*" published in 1979 was a standard work on the subject (Published by the IEE ISBN 0 906048 04 4)

The 37900A/B Signaling Test Set

This is good place to address a conundrum – the spelling of "signalling" or "signaling". HP, being an American company, meant the US spelling (or mis-spelling) prevailed on the instruments and literature even though it was all produced in Scotland. In the remainder of this chapter I will continue to use the UK spelling except when referring to an instrument by its "proper name" as it appears on the front panel.

The first product was the 37900A/B Signaling Test Set which launched in the 1989 HP Catalogue, published at the end of 1988 only 12 months after the Telenorma deal. The Telenorma signalling processor cards (effectively implementing the Message Transfer Part of the protocol but without the physical interfaces) were repackaged and installed in the HP 9000 Series 300 card-cage chassis. Various physical interfaces were used in the network so these had to be defined and designed and offered as options for the signalling processors. David Guest, the project manager, recalled that this development needed to be completed in a very short time and this was undoubtedly helped by regular morning meetings when difficulties and solutions could be shared between team members.

The signalling processor part of the system had the model number 37901A, and the strategy was to mirror developments at HP's computer divisions in Fort Collins and Corvallis with a view to future networking opportunities. For the 37900A version, the main software and user interface was built into the new HP ES/12 Vectra PC (HP's first IBM compatible PC). The "A" version was intended as the entry model with monitoring of one bi-directional signalling link, whereas the "B" model would have a higher specification with monitoring of four signalling links.

The Telenorma software ran in Pascal language, so the Vectra PC was fitted with HP's Pascal Workstation on a card. The attempt was to provide a lower-cost entry to the signalling market no doubt to appear competitive with established protocol analyzers, however the Pascal card developed in the US was a poor implementation. Graham Byars, who was involved on the marketing side, remembers the production assembly was in a room on the ground floor at the factory.

"*It took about ten PCs to get one that actually worked, and the air was generally "blue" in that room, but it was a great learning curve on what needed to be fixed.*"

Peter Locke made a similar observation:

"*We always planned to release the product on the HP Series 300 Pascal Workstation, but the introduction of PC-based product was an idea late in the design cycle to build up interest in the market by providing a low-cost entry. Once we were aware of the problems, we pulled it as quickly as possible, however it provided some useful market exposure and demonstrated the necessity of multi-link operation provided in the "B" model. We got roasted by the European customers when we took the single-link solution on the road!*"

The Division did sell one 37900A at the end of 1988, however it was sold to Royal Air Force. The consensus was that the RAF didn't have SS7 so this appeared to be a way for them to acquire a Vectra PC by the back door. Certainly they never accepted the free upgrade when the signalling tester was fixed!

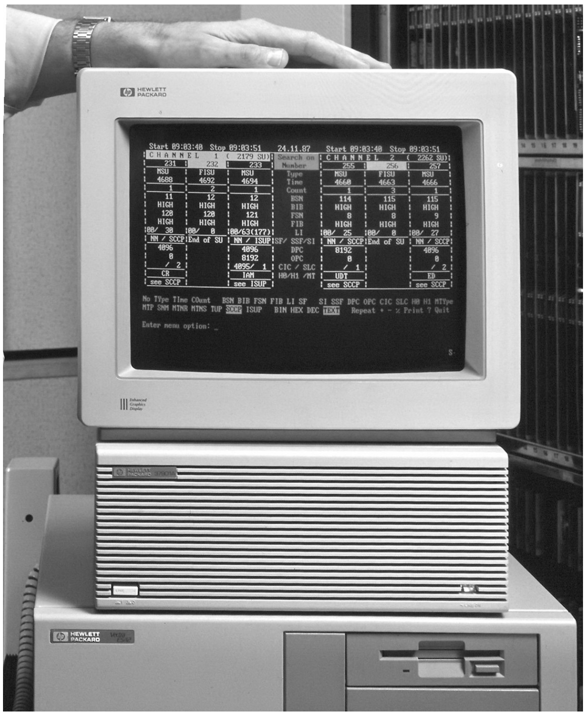

This photo shows the 37900A at introduction in August 1988, exactly a year after the Telenorma deal was signed. The display shows two signalling links being monitored using the Telenorma software running on the Pascal card in the Vectra ES/12 at the bottom of the stack. In the middle is the 37901A signalling processor which could handle two bi-directional signalling links.

However, the ineffective "A" model was quickly dropped in favour of the 37900B which used a proper HP 9000 R/332 Pascal workstation as the development team had originally intended. This added to the number of boxes, so in addition to one or more signalling processor units, there was also the computer, possibly a disc drive, a terminal and a keyboard. Quite a stack of equipment – suitable for the lab or production, but unattractive for field installation. It was also quite a handful for the marketing staff to take on demonstration tours! Apart from its physical size, this system did the job and was more powerful than any of the standard protocol analyzers on the market, as discussed earlier. It could control two signalling processor units and so could monitor four bi-directional signalling links simultaneously.

The 37900B was introduced in the summer of 1989 and was moderately successful considering the early stage of the SS7 market. The market numbers are a bit uncertain as various configurations were possible, but records show that 150 units of 37901A signalling processors were sold with the 37900B which might equate to around 100 systems assuming a mix of two- and four-link configurations.

The 37901A signalling processor contained multiple processor cards which interconnected with a range of optional signalling link interface cards. These handled primary rate (2.048 Mb/s and 1.544 Mb/s), 64 kb/s (DS-0 and V.35) as well as RS-232 (and RS-449). The most popular ones were the primary rate options[10].

As well as monitoring up to four signalling links, the 37900B could be configured to generate SS7 traffic. This was called emulation, and the computer system was programmed so that the user could define the signalling traffic using CCITT's Specification and Description Language (SDL). This is a high-level language making specific commands more explicit, a further differentiator from standard protocol analyzers.

37900C Signalling Test Set

Around this time, HP computer divisions in Fort Collins and Corvallis introduced an integrated version of the R/332 workstation in a portable package. It used a smaller monochrome CRT screen and a fold-down keyboard. There was space for plug-in cards at the back, and one of the design team, Peter Locke, remembers working with the US

[10] The total number of interface cards sold with the 37900B/C were 37911A (2 Mb/s) 200 units; 37912A (1.5 Mb/s) 80 units; 37913A (RS-449/232) 30 units; 37914A (V.35/64 kb/s) 96 units.

divisions to ensure the unit would be suitable for South Queensferry. Since this was simply a repackaging of the computer hardware, the same software and signalling processor and interface cards could be used as in the 37900B. The main difference was the space limitations in the portable box meant only two links could be monitored compared to the maximum of four in the "B" model.

The computer group had codenamed this workstation "*Rainbow*". Graham Byars commented that perhaps the Division hoped the 37900C would lead to a pot of gold at the end of SS7 rainbow. While it didn't do that, it did produce reasonable sales as around 170 signalling processors were shipped with the unit. Although it was an all-in-one test set, it was still a heavyweight package, nearly 28" (70 cm) long and weighing about 50 lbs.

Dave Warren, one of the project managers, commented:

"*The 37900B/C were unique products because of their high-level decodes, triggers and filters, but still had awkward packaging (multiple heavy boxes linked by heavy cables) and very much a computer application rather than an instrument (you "booted" into Pascal Workstation then ran your application). They were already well respected in the R&D market but we realised that they were not practical for network operators to use.*"

Both the B and C versions were launched in August 1989 and were in production until the summer 1991 when they were replaced by the 37900D. With this new product, South Queensferry did at last find "the pot of gold at the end of the rainbow" – in fact there were several pots of gold!

37900D Signaling Test Set (Unit Preserved NMS T.2010.70)

In the late 1980s, the Division embarked on a couple of new investigations aimed at using the latest programmable logic, signal processing and software developments.

One was called ALBERT and would create a highly flexible pattern generation and detection product with all kinds of framing and other higher-level functionality. It would target new datacommunications applications and emerging markets such as ISDN. It had a top rate of about 20 Mb/s, and planned to use the VXI card-cage system.

The other investigation was called VICTORIA and would use a similar approach to create the next generation of signalling testers and systems, moving on from the Telenorma-based products. In the 1990 strategic plan for the next five years, the signalling product plan showed how the 37900B/C would be replaced by VICTORIA, as illustrated here.

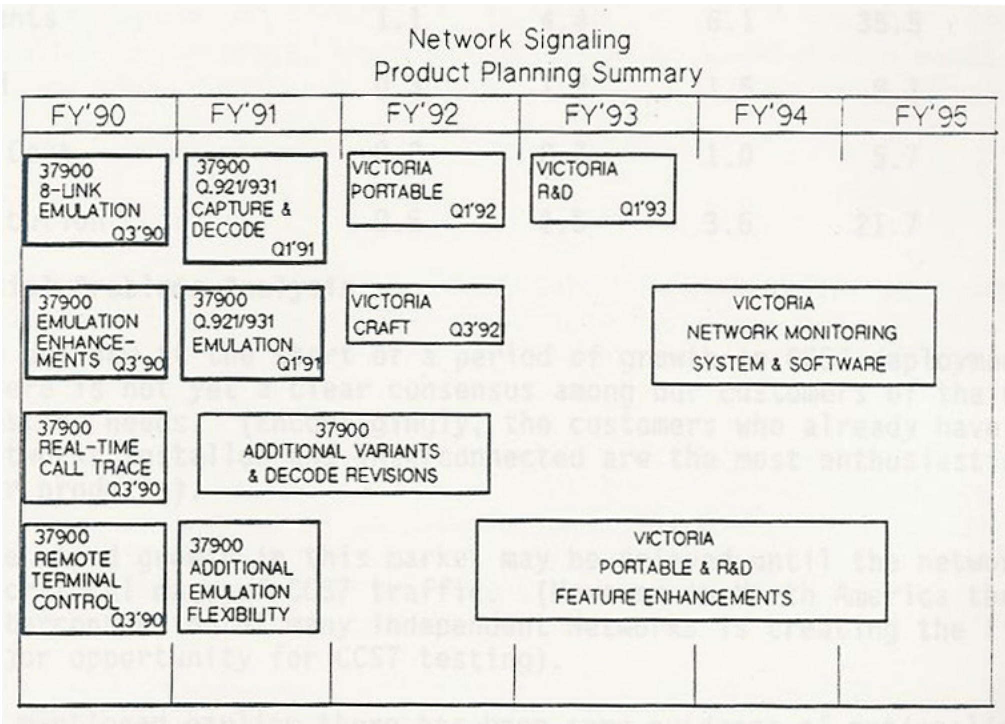

There was still plenty of on-going work on the 37900 and this was headed by Ian Burrows, while the new VICTORIA project was led by Dave Warren. In a project report from early 1990, Dave Warren commented that a total of six engineers were working on VICTORIA and they were developing close links with Colorado Telecommunications Division with a view to using some of their "Pinecone" protocol analyzer platform. What Pinecone was, is now somewhat lost in the mists of time.

However, sometime in the following 12 months, there was significant change of plan. The ALBERT project was abandoned (see Appendix 6), because the market was moving rapidly to higher-speed SONET/SDH, and Queensferry had acquired the necessary chip technology from Alcatel (see Chapter 10). The strategy on VICTORIA also seemed to change. David Dack, who as R&D section manager had led the group into the signalling field, left South Queensferry to join the new HP Labs operation in Bristol in May 1989, leaving his successor, David Guest, with the decision on how best to proceed.

The signalling market was evolving rapidly by the early 1990s, with suppliers of protocol analyzers entering the market. The main competitor, Tekelec Inc., had recently acquired a company called Protocol Technologies who had a traffic generating product, MGTS (Message Generator Traffic Simulator), and there were new products from Siemens and Elmi. The 37900B and C did the job well and were established with several customers who used them for conformance testing of new SS7 systems[11]. However, the products were beginning to look quite cumbersome, and the VICTORIA project was still at least two years away from delivering products. Furthermore, division management was becoming sceptical about the lengthening timescales and potential market. A faster route was clearly called for.

In the late 1980s, the digital transmission team had started building a new generation of portable testers into a package codenamed "Hornet", as described in Chapter 10. This borrowed the mechanics from the highly successful 8590 series of spectrum analyzers, some of which were produced in high volume in the neighbouring Queensferry Microwave Division. The Hornet box was attractive as it came with all the necessary hardware including protective rubber feet, a front carrying handle and a strong fibreglass cover with latches to protect the front panel. It was an ideal platform for the field maintenance and installation market. The question probably occurred to the signalling group: could the next generation test set also be built into a Hornet package?

The idea was born of getting to market much faster by re-engineering the Telenorma design into the Hornet package, and this would become the new 37900D Signaling Test Set. This would also capitalise on the objectives and development work already done on VICTORIA, including a single portable box and four-link monitoring. Peter Locke's team worked on a greatly improved call-trace facility, while Ian Burrows worked on condensing the interface cards, and Dave Warren on the 37900D chassis and user interface.

Call Trace was a new measurement which used the powerful filters and real-time decoding in the 37900 to identify all the signalling messages associated with a particular telephone number which was encoded in the higher-level User Part. From this, the instrument derived the Call Data Record. Dave Warren remembers a demo they did to Finlay Mackenzie and to engage his interest, they used his home phone number and pretended his wife had been trying to call him, but couldn't get through. "*From that moment he was hooked!*"

Ideas for the new signalling tester in the Hornet package were coming together. The package was adapted to take five modules plugged-in at the back. HP had condensed the Rainbow Pascal processor onto a single card and this occupied the bottom slot. The four slots above took up to four interface/signalling processor cards. Previously, the interface and signalling processor cards had been separate, whereas the new design combined them. The new signalling cards had two receive ports able to monitor a bi-directional signalling link and two transmit ports. A fully loaded test set could therefore monitor four bi-directional signalling links (a typical "link set") or emulate up to eight links simultaneously using the transmit ports.

---

[11] Bellcore Labs (the research arm for the Bell Operating Companies) built their SS7 training material round the 37900, so this was a valuable reference when selling the test sets in the USA for years after.

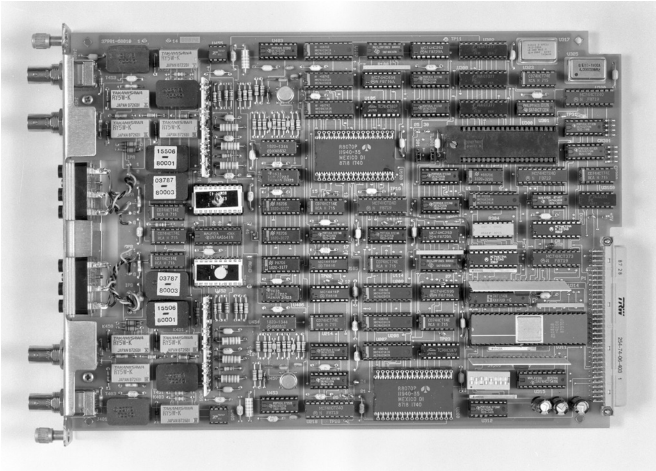

Work on combining the interface and signalling processor cards was a real challenge as it required doubling the component density on the boards. The multiple connectors were quite large and the interfaces had to withstand high voltages. They were some of the highest-density boards produced in South Queensferry at the time, an example shown here.

Dave Warren commented on the challenging design of the 37900D chassis:

"*Firstly, it had to ship in less than a year after definition release, whereas typically it took two to three years. Secondly, it had to include a computer, four large interface cards, a hard disc drive, a flexible disc drive, a large display, a full QWERTY keyboard and a 250W power supply in a single portable package. This was the era before laptop computers, so small computer components did not exist. The only way we could meet the aggressive timescales was to reduce the prototype cycles from typically five down to two. Our product designers, working on the metalwork, were the first to use 3D modelling.*"

The design team had the clever idea of getting a specially made keyboard to fit inside the Hornet front cover which was then hinged to the bottom of the front panel. Most of the instrument front panel was occupied by a large (9" diagonal) electro-luminescent display. It was a compact instrument, ideal for field use. With all its protective rubber feet etc., it was just under 19" (480 mm) long and 14.5" (370 mm) wide and weighed 13 kg.

Peter Locke commented that signalling needed a new perspective from the management team. Previously, most products had been predominantly hardware-based, although incorporating an increasing amount of firmware, and with these there was a definite end-point to the development when the product transferred to production. In the signalling market there was a need for regular updates on software. As Peter recalled:

"*I had set up a rolling programme to release new features every six months which included customers in the decision process. Customers loved this as signalling was a very fast moving area and we were the only ones committing those kinds of resources.*"

Graham Byars, Marketing Engineer on signalling test, made a similar observation:

"Customers loved the open nature of our software. It was structured so modifications and enhancements were possible, whereas our competitors had closed architectures which tied the customer into high on-going software costs."

Nevertheless, the project did have one or two difficult moments as Dave Warren recalled:

"*At one point in the Development cycle we knew the future of the project was in doubt, however Graham Byars was still marketing it to the leading Network Operators. I remember a Friday morning when Graham was at Swisscom and rang to tell me we had* *just won our first big order for 50 units, if we could ship them in four months. About an hour later Finlay called me to his office where he and Robin Myles told me that they couldn't afford to continue the product. I replied 'What do you want me to do about this Swisscom order for 50 units?', (they hadn't heard yet). They both smiled, and Finlay asked if we could ship them that financial year. We did!*"

What a happy coincidence that the Swisscom order came along when it did, otherwise the history of HP South Queensferry might have been very different!

Dave Warren also remembered a marketing slogan used to promote the 37900D to the sales engineers: "*Half the size, half the weight, and half the price but with twice the number of interfaces*". It was certainly a big step forward from the earlier models. The development was completed remarkably quickly in under 12 months and the new 37900D entered the market in the summer of 1991. It was a brilliant piece of design. The team was "firing on all cylinders".

The 37900D took off, and shipped 200 units with a revenue of $12M in the first three months, as much as the previous units had managed in two years. The marketing staff came up with an excellent four-page colour sales brochure that emphasised the advantages and benefits of the product without excessive technical detail. *"Unlock the secrets of network signaling"*, and *"Follow events as they happen, understand signaling interactions and trace message sequences to validate call set-up – the easy way"*. It would have been a great tool for the sales force to introduce the product. A second brochure followed, more focussed on the mobile telephone market, a rapidly growing sector in the early 1990s.

The advent of mobile phone networks immensely increased the complexity of signalling. As mentioned earlier, SS7 has an important role in the mobile network since, when a call is set up, various transactions are necessary between the Home Location Register (HLR, the database holding the subscribers registration), the Visitor Location Register (VLR, which logs the current position of the mobile phone as it roams around) and the Mobile Switching Centre (MSC) which routes the call. These rely on some specific higher-level Application Parts and User Parts in SS7. These were available as an option on the 37900D, and indeed the flexibility of the instrument software allowed many variants to be produced including local language versions. One variant recalled by Ron McDowall (marketing) was a version for the Japanese cellular market including a special 2 Mb/s interface and SS7 protocol stack, which was done with a small design team in Kobe, Japan. "*Motorola bought a good number as they were supplying base stations to some new mobile service providers as the wireless market in Japan began to open up.*"

Graham Byars recalled some of the key customers in the mobile sector from those "*Heady Days*" of the early 1990s. PKI (Lucent) in Nuremburg who were building the first GSM base stations bought about 200 units, and Alcatel SEL also in Germany who were building core Mobile Switching Centres (MSC), HLR, VLR etc., bought over 200 units. In the USA, Motorola was a huge customer buying around 300 units. The mobile market proved to be another "pot of gold" for the 37900D.

With these customers and others, South Queensferry built a close relationship so that software could be modified to their requirements and in some cases they had access to the source Pascal code to write their own routines for the instrument.

The 37900D was in production from 1991 to 1999 and it is thought that the Division shipped around 2000 units worth $170M, although definitive sales figures no longer exist. We do have specific sales statistics for the five years 1992 to 1996, which give an idea of its contribution:

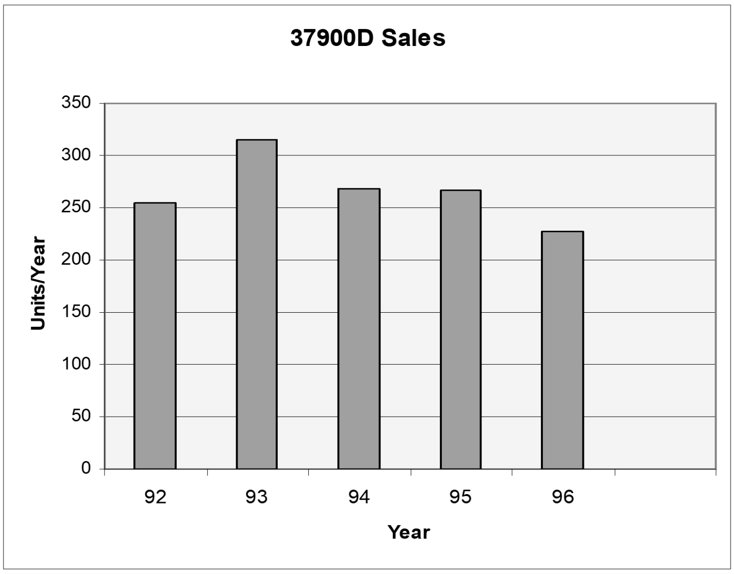

The signalling processor/interface cards that plugged into the instrument were also in demand as separate items, used by the HP Computer Group to build servers and databases for Service Control Point (SCP) applications in fixed and mobile networks. Peter Locke recalled some initial work done jointly with HP Labs in Bristol which used South Queensferry's signalling cards in a UNIX system. Following on from this, Telecom Networks Division in Grenoble, France developed the HP OpenCall SS7 Platform[12] using the South Queensferry hardware.

The 37900D had the power to drill-down to a specific message sequence for an individual call, so it could be used to identify a particular telephone number and the calls that were being made. Sometimes the unit was used for "intelligence gathering", although who or what was under surveillance the factory was not supposed to know about. As well as these detailed searches, the instrument could go all the way up to making a high-level analysis of overall traffic statistics and generating multiple SS7 traffic streams for system testing.

The 37900D preserved in the National Museum collection is an early unit manufactured in mid 1991 and the 204th of the overall production run.

[12] HP Journal, August 1997, http://www.hpl.hp.com/hpjournal/97aug/aug97a7.pdf

During the 1990s, the SS7 market expanded rapidly, particularly in the mobile sector. In the Division's 1990 strategic plan, the market was forecast to grow at 25% per annum, which begs the question why the 37900D has the flat sales profile shown above. One reason was probably that Siemens brought out a new SS7 test set[13] with a Microsoft Windows user interface. This became popular with PC users in the 1990s (initially Windows 3.1 followed by Windows 95). The 37900D script-based screen display looked more like something from the 1980s, which indeed it was. For more technical users in R&D, manufacturing and installation, this wasn't much of an issue, and they may well have preferred it. But for service providers wishing to monitor network performance and troubleshoot problems, the Siemens Windows interface was preferred for ease of use. Graham Byars conceded that the 37900D lost market share, although the emulation business held up well.

So, why not a revamp of the 37900D with a Windows user interface? Well, for the moment the signalling team had some "bigger fish to fry".

A Systems Business

In the long history of HP, perhaps only a handful of products can be said to have launched a new division, but the 37900D is probably one of them. In the 1990 product plan referenced earlier (page 204), there was a "VICTORIA Network Monitoring System and Software" shown for 1994/95. VICTORIA had gone, but the idea lived on, and as the SS7 network expanded and South Queensferry's knowledge and experience of the market increased, it was looking like a good opportunity. The Division already had experience of network monitoring and remote test having introduced FDM monitoring systems in the late 1970s, and later the computer-controlled RATES systems used for checking analogue private lines (Chapters 5 and 7).

In fact South Queensferry had developed considerable expertise in the early 1980s in the design of large software systems for network management, usually called Operational Support Systems (OSS) in the telecom industry. The Division's platform was called TSS/1000 (Telecom System Software running on HP 1000 computer). As mentioned earlier, David Dack was a key driver in this along with Robin Myles who directed a number of software engineers in the Software Products Group which focussed on these systems. So, it was not surprising that early investigations into signalling test also included ideas for monitoring systems.

By the early 1990s, SS7 installations had reached the point where operators were becoming aware of the large-scale traffic management problems of running a network of computer-controlled telephone switches. Imbalances caused by failures, and surges of traffic caused by mass-calling events, could lead to network instability and blocking, even shutdown. Also, in the deregulated service provider market which evolved from the mid-1980s onwards, along with the growth of mobile phone networks, there was the need to exchange traffic across network boundaries with call set-up and billing information, which was all handled by SS7. There was no way that portable signalling test sets would ever be able to keep track of all this real time information and analyse it on hundreds of signalling links.

[13] Initially the Siemens K1103 and later the K1197 which had a very similar specification to the 37900D even down to its size and weight. Siemens test sets were eventually bought out by Tektronix.

HP acceSS7 Monitoring System

Around this time work started in earnest on defining a network monitoring and management system for SS7, initially using the technology and software developed for the 37900D. This new system had the title "HP acceSS7" and the product number E4250A. The Signalling Group's resources were now focussed on this new opportunity. The Division's Marketing Manager at the time was Tom White who had previously been a sales engineer and sales manager in the UK, and was well aware of the systems opportunity through the BT contracts for RATES. Armed with the sales success of the 37900D and developments in the market, he put a proposal to his corporate bosses in Palo Alto outlining a plan for a new division focussed on systems and especially SS7 monitoring, separating it from the QTD instrument business. Eventually, he got backing of several million dollars to set up a new division at South Queensferry in 1994 called the Telecom Systems Division (TSD). Tom White also became the Division Manager, a factor which no doubt drove his motivation!

As mentioned, many of the initial ideas for the new HP acceSS7 came from simply scaling-up the hardware and software in the 37900D Signaling Test Set. Whereas the test set could only monitor a maximum of four bi-directional signalling links simultaneously (the typical link-set connection between two signalling nodes), the new system would be capable of monitoring hundreds or thousands of signalling links across the network simultaneously, and so could give a real-time view of how the whole network was behaving. No stand-alone instrument could ever achieve this integrated view.

The basic architecture of the acceSS7 system is shown in the diagram (as depicted in the 1996 HP Catalogue). The remote site consisted of measurement card-cages which contained multiple monitoring cards which accessed the signalling links non-intrusively. In principle, these were similar to the 37900D interface and signalling processor cards, though of course no transmit emulation capability was needed. The measurement cards could be scaled up to monitor a large number of signalling links simultaneously, and these interfaced to a remote site processor which aggregated and condensed the large amount of information being collected.

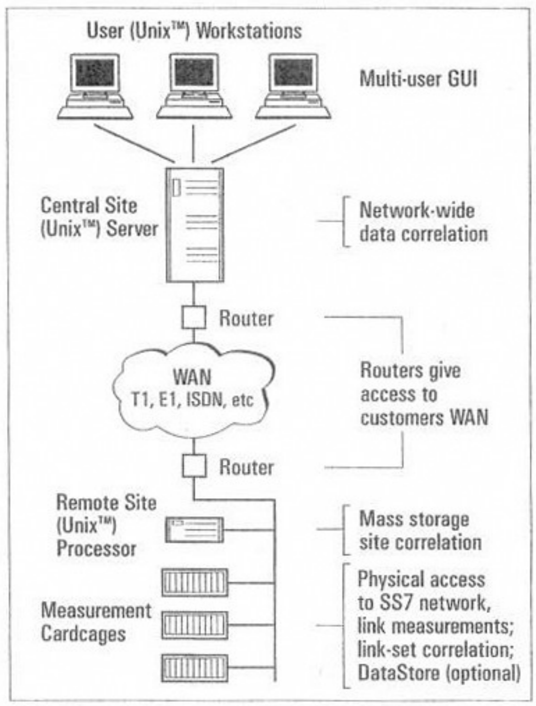

Structure of HP E4250A system

Multiple remote sites could be connected via the Wide Area Network (e.g. 2 Mb/s digital links) to the Central Site Server which ran the main acceSS7 software providing the network-wide correlation. User terminals (Unix Workstations) were connected to the central server, and users interacted with the system through a menu of "windows" showing SS7 transaction data and traffic statistics. The graphics displays of statistics were a quick way of identifying abnormal traffic loads and calling patterns, as shown in this example, allowing the operator to see the big picture and then focus in on an identified problem:

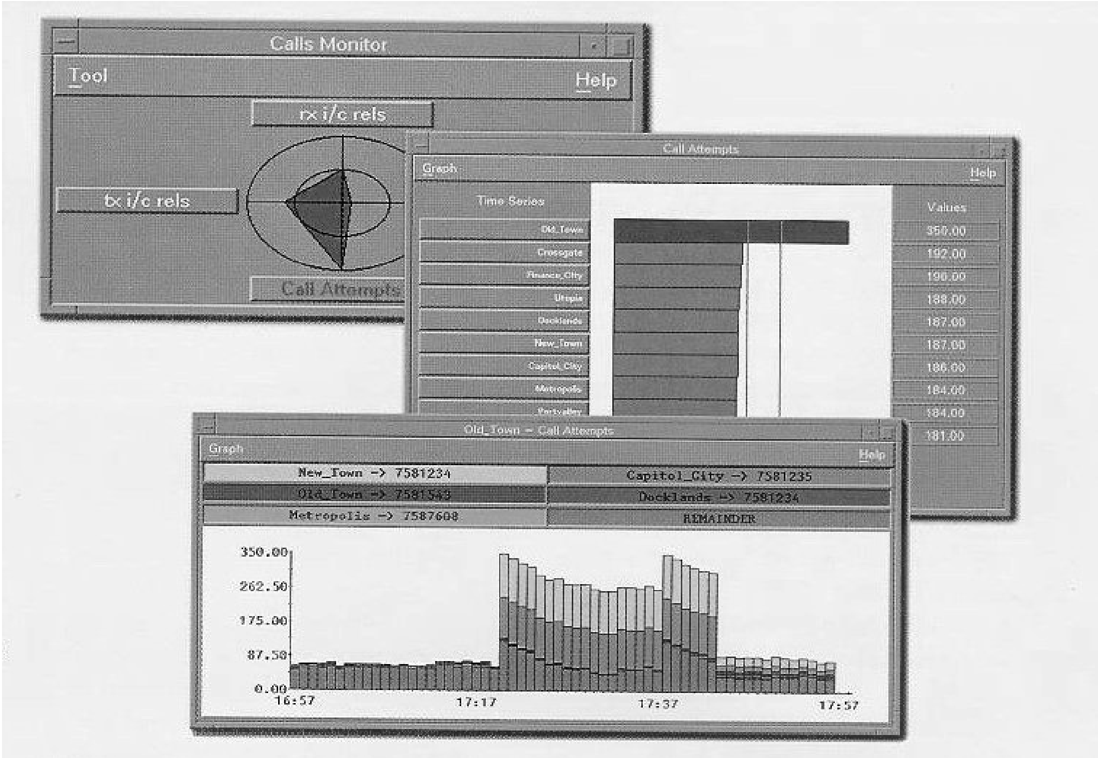

In contrast to the monitoring embedded in the operational equipment, HP acceSS7 claimed to give network operators a truly real-time network-wide view, identifying degraded system performance and imbalanced traffic loads that might give early warning of potential problems.

Apart from software development, design work was needed on the monitoring hardware which was to some extent based on the signalling processor cards from the 37900D. These photos show an early monitoring rack and a bank of 37900D Signaling Test Sets used to provide a traffic load for testing the monitoring system. A couple of the monitoring cards are shown on the next page.

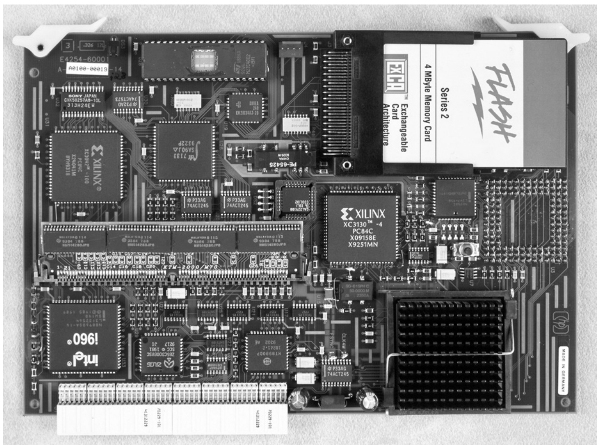

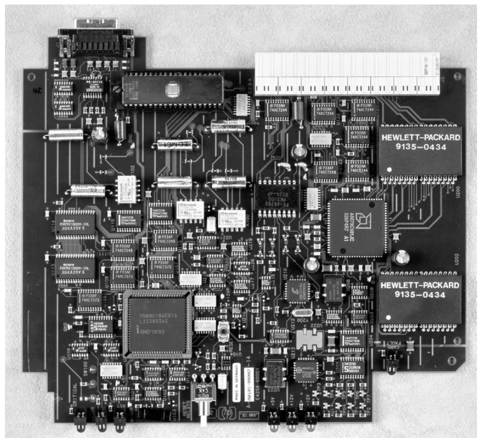

Two types of data were derived from the HP acceSS7 system: firstly traffic statistics and network status, and secondly, traffic content.

The traffic statistics were useful to monitor network efficiency and identify under-used and overloaded sections of the network including switch capacity. This was valuable for network provisioning and for evaluating quality of service to customers.

The second class of measurements were more focussed on managing the business and relied on the capability of identifying and tracing a particular call or group of calls from a particular telephone number or to a particular number. This information could be aggregated over time in the Call Detail Record (CDR) for later analysis, sometimes called "Data Mining". This was a development of the Call Trace capability in the 37900D that had been a key selling point with the stand-alone instrument.

The Call Trace feature had quite a lot of different applications. It could be used to troubleshoot incorrect call routing and cases where there was a number translation initiated through a Service Control Point (SCP) database, such as call-forwarding or Freephone 800 service. This gained in importance as network services became more complicated and there was more traffic handover between independent networks, for example between mobile and fixed networks.

By the late 1990s, TSD was developing a number of important new applications using this concept, to help network operators manage and protect their business. One application was gathering Billing Data for traffic transiting one network to another. Another very popular application was Fraud Detection. It was thought that service providers could be losing billions of dollars a year to various fraudulent calling scams. Unless the culprit could be identified quickly, it would be difficult to catch up with them and prosecute. The acceSS7 fraud application allowed service providers to identify abnormal calling patterns very quickly and shut the fraudulent operation down. Through the "open system" architecture of acceSS7, TSD worked in partnership with other software developers in fraud detection to offer solutions in a rapidly evolving market. There is more information about this business in Volume 1 Appendix 2.

These powerful call trace applications, and aggregation of call data in the CDR, also became of great interest to security agencies such as GCHQ and the National Security Agency (NSA) in the USA, particularly after the 9/11 terrorist attack in 2001. Sometimes referred to as "Meta Data", this signalling information can yield a lot of detail about who is calling whom and when, their possible motivation, the location of mobiles and so on, without ever listening to the content of the call itself.

By the early 2000s, TSD was making most of its revenue from these new business applications, in a way transcending the original objectives of monitoring network performance. It was a success story, with revenue peaking at over $200M, much of it from these business applications. This success was recognised in April 1999 when TSD received the Queen's Award for Technological Achievement.

I have covered the business history and fortunes of TSD in more detail in Chapters 10, 11 and 12 of Volume 1, however, this whole activity would probably never have happened without the 37900D. Indeed the 37900D continued to play a role as it helped fund the venture in the early days. Graham Byars recalled the marketing team being given $500k per year to bring in $10M to $12M of sales – in effect the 37900D was a "cash cow". Ron McDowall, who was Product Manager at the time, remembered they were given the challenge of reviving flagging sales to help fund the development of HP acceSS7:

"*A small R&D team headed by Ian Burrows 'repackaged' the product as a basic platform with a set of software measurement modules which we re-priced and remarketed. I think it may have been the first time there was separate pricing for some of the software on the box. The sales did increase to over $10M, and so generated some income in the early years of TSD*."

Before the new "acceSS7" system was ready, customer interest was also kept alive with some system solutions using remotely controlled 37900D test sets, in which the display and control of the remote instrument was replicated at the controlling computer. Peter Locke recalled some innovative work by John Duff, and also Bill Lauchlan in Canada who had previously worked on monitoring software for analogue FDM systems. Later the 37900D was an essential test tool in evaluating the HP acceSS7 system when banks of instruments were used to generate an SS7 traffic load.

So, in a way the 37900D created another "pot of gold" at the end of the SS7 rainbow.

37907A Signaling Advisor (Unit Preserved)

With much of the R&D effort in signalling test now applied to the new opportunity in monitoring systems, further software work on the 37900D was undertaken by a subcontractor, Libra (Syscom) Ltd. Libra was formed in 1994 and occupied one of several small units created on the ground floor of the South Queensferry factory in collaboration with Scottish Enterprise. This was known as the Scottish Software Partners Centre and was intended as an incubation facility for start-up businesses, and thereby put some spare factory space to productive use. The lab team established a good working relationship with Libra which continued through the second half of the 1990s as Libra expanded. Being just downstairs from the main lab, it was very much like an extension of the R&D team.

In 1994/95, the Colorado Telecom Division (CTD) introduced a new protocol analyzer platform, the J2300 Internet Advisor series. This was like an early laptop computer with a colour LCD display and a Windows-type user interface which hinged-up to access the keyboard. It was "chunkier" than a modern laptop as it incorporated various optional physical interfaces for telecom and LAN applications.

South Queensferry decided to use this as the platform for the next SS7 test set using Windows 95. This became the 37907A Signaling Advisor (initially it had the model number 37900E). To provide the additional physical interfaces and signalling processors for multi-link applications, an optional under-cradle accommodated the extra hardware. Although the project was managed by Dave Warren in the R&D lab, much of the work on both hardware and software was undertaken by Libra, which hired more engineers to fulfil its contract with HP. The different physical interfaces required to monitor signalling links were provided by a range of interface modules, which plugged into the side of the Advisor. This photo shows the E7572A module for 2 Mb/s signalling links. This module can monitor four bi-directional signalling links, illustrating how new miniaturised components had increased the functional density compared to the earlier 37900D interfaces.

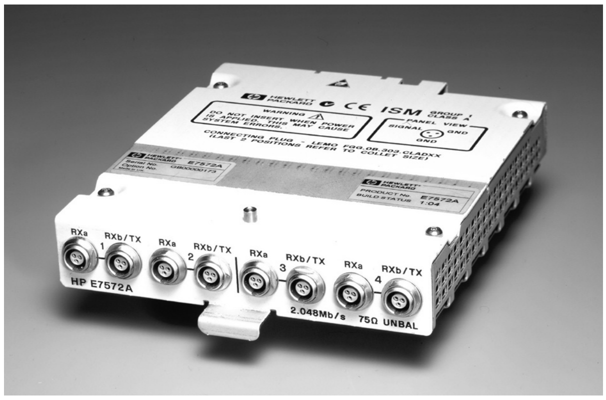

The diversity and expansion of SS7 meant there were many more detailed variants and extensions than expected, so it was necessary to provide a continuing flow of decode software. Libra had underestimated the amount of work needed to complete the Signaling Advisor, and had to be bought out by HP in 1999. The 37907A (pictured here with Neil McKinlay from marketing) was finally launched that year, and the 37900D discontinued.

David Guest recalled that around this time HP was increasing its presence in China, and engineers were hired there to provide additional effort in software development for the signalling business.

"*Training visits to Queensferry for the Chinese engineers was an enjoyable experience for those at both ends of the relationship.*"

Shortly after this, there was another change of direction.

In 1999 and 2000, South Queensferry experienced an enormous boom in sales of digital transmission test sets (an aspect of the worldwide "Dotcom Bubble" that ultimately proved fatal to the Division). Signalling test was no longer part of the core business and the Division was encouraged by corporate management to focus exclusively on the transmission test market. The obvious solution was to transfer the Signalling Advisor to CTD where it would naturally form part of the Internet Advisor family. By then the SS7 systems business was an independent division, so was unaffected by these changes.

Around 2000/2001 the Signaling Advisor was transferred to CTD and about the same time, HP's test and measurement business became Agilent Technologies, so the product was re-branded the Agilent Signaling Advisor. In the 2001 Agilent catalogue, three versions of the 37907A were marketed: a version for SS7 (J4211A) and two for cellular mobile networks (J4212A for GSM and J4213A for North America). As with the 37900D, a range of physical interface modules were available, which plugged into the side of the Advisor.

The final twist in this story of signalling test, takes us forward to 2010. In the ten years that followed the Signaling Advisor's transfer to CTD, many changes took place in the Agilent business portfolio. The South Queensferry SS7 systems business at TSD grew and then contracted after 2007, possibly as a result of poor management and a lack of strategic interest from the Agilent Corporation. The full story of this business is told in Volume 1 Chapter 12.

In early 2010, as part of its realignment, Agilent sold both the SS7 systems business and the protocol analyzer business in Colorado for $160M cash to another test equipment vendor, JDSU, a company that had previously absorbed Acterna, itself a merger with South Queensferry's main competitor over the years, Wandel & Goltermann.

And so came to an end one of the most interesting of the product stories from South Queensferry. It also marked the end of the telecom test business at the South Queensferry site.

Conclusion

Let's round this chapter off with a brief recap. A couple of engineers get interested in signalling and by chance are introduced to an old guy who turns out to be a world authority. Coincidentally, a German company offers to sell them a signalling test set development. The early sales are not very encouraging, but they persevere and the third product is a notable success, so good in fact, it leads to the launch of a new division which, a few years later, wins the Queen's Award for Technological Achievement for its monitoring system. The product evolves far beyond the original idea of checking message integrity and becomes a tool for running the telecom business, managing traffic, detecting fraud, tracking billing, and even assisting national security. This was surely a business with a future …. if the vision that started the enterprise in the late 1980s had still been there.

Acknowledgements

The following former employees at South Queensferry have contributed to this chapter: Graham Byars, Peter Locke, Finlay Mackenzie, Robin Myles, Dave Warren, Ron McDowall, John Duff, David Guest, Boyd Williamson and Ivan Young.

12

Chapter Twelve

Bus Extenders and Modems

This chapter describes some interesting products, developed at South Queensferry in the 1970s and 80s, that had only a tenuous connection to the Division's mainstream charter of telecom test. However, by "straying" from its charter, the Division tapped a rich seam of general-purpose applications way beyond the communications business, and some of the products were real "money-spinners" for the factory, generating substantial profits for many years. This is the story of how these unique products came about and why they were so successful.

Hewlett-Packard Interface Bus (HP-IB)

In the mid-1960s, HP became interested in instrument systems controlled by computers and acquired some embryonic companies[1] involved in this emerging area. While Hewlett and Packard could see the emerging market in scientific computers, they decided to focus particularly on the area closest to their business, which was automating measurements using an instrumentation computer. The first HP product, the 2116A computer was introduced in 1966, and it differentiated itself from other offerings by its robust construction (similar to HP instruments) and its architecture which focussed on a multi-channel input/output (I/O) structure for simultaneous control of up to 48 instruments. It was also one of the first computers to use integrated logic circuits rather than discrete component logic, which made it a compact unit for its time, albeit weighing over 100 kg!

To complement the 2116A, HP offered an expanding range of instruments with a remote control interface, allowing the computer to control measurements and receive the measured result from the instruments. The interface on the computer and the instruments was usually a parallel data interface, sometimes separate connectors for input and output. For

[1] Specifically, Dymec and Sanborn who were involved in instrument systems and data recording.

example the well-known 5245L frequency counter had two 36-pin "Amphenol" connectors. One of South Queensferry's first products, the 3722A Noise Generator had a single 36-way connector for remote programming. Separate physical interface cards had to be plugged into the 2116A computer for each instrument being controlled via the 36-way interface, and the computer bus polled the interface cards and operated a priority interrupt system to ensure efficient operation. The control software tended to be specific to a particular instrument as the interface commands and data structures weren't standardised.

The Company built some effective automated measurement systems using this architecture, and in the early 1970s South Queensferry set up a Systems Integration Centre to bring together products from various divisions, including Scotland, under computer control along with the custom applications software.

The growth and success of this business led HP to consider how the instrument control architecture could be improved to make it more flexible and less cumbersome than the multi-port parallel I/O structure of the 2116A computer. Conceptually, the inventive step was to take the bus structure out of the computer backplane and integrate it into the instrumentation and interconnections. The computer then only needed one I/O port, forming part of this external bus structure. Thus was born the Hewlett-Packard Interface Bus or HP-IB, soon to be adopted by the wider world.

When the HP-IB system was developed in 1971-72, the designers had to find a compromise between speed and practicality. A completely serial interface (for example the Ethernet we have today) would have been simple physically, but too slow and complicated with the available electronics at the time, whereas a fully parallel interface with every function and data field having its own separate connection, would have resulted in thick cables and cumbersome connectors. The HP-IB is a compromise (bit-parallel, byte serial), with an 8-bit data bus used for transmitting information and addressing, plus eight control lines, making 16 connections in all. The three engineers associated with this development were Gerry Nelson in Loveland Division, Dave Ricci in Santa Clara, and Don Loughry at Corporate who had previously been with the Dymec Division[2].

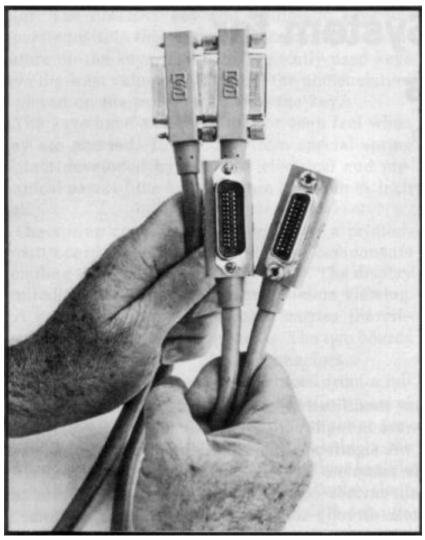

An important part of the standard is the HP-IB connecting cables which have a male/female back-to-back connector at each end, thereby enabling them to be stacked (as in the photo). This allows instruments to be "paralleled-up" or "daisy-chained" with up to 31 devices on the bus and a maximum cable length of 20 metres, which is more than enough for most racked test systems or bench measurement set-ups.

There are 24 conductors in the HP-IB cable, but as mentioned earlier there are 16 active connections. The diagram below shows the key features of the bus. The devices connected in parallel across the bus can be controllers (for example a computer or desktop calculator) and various other devices such as peripherals or instruments that are talkers or

[2] A good record of this early work done by these designers, and the principles of HP-IB, were published in two articles in the HP Journal of October 1972: http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/1972-10.pdf

listeners or more usually both. The electrical logic signals on the bus are "ground-true" and the voltages are those of TTL (Transistor Transistor Logic); that is nominally +5V and 0V.

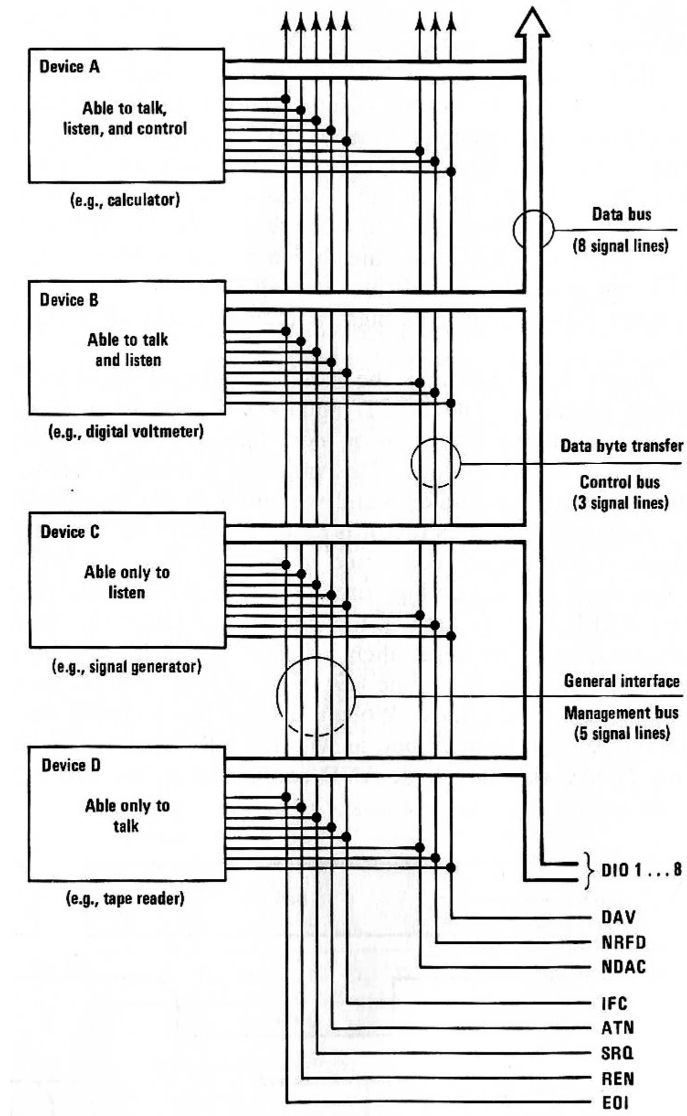

One or two facets of HP-IB operation are worth covering here as they have a bearing on the design of the South Queensferry products described later. Firstly, the 8-bit wide data bus is used either for transmitting parametric data to and from devices connected to the bus, or for sending address and control information. This alternative functionality is controlled by the ATN (Attention) line which is normally high but when it is set low ("true") by the system controller all the connected devices read the bus as control/address information instead. Each device has a unique address and when this address is present on the bus, that particular device will be activated to send/receive data until unaddressed later. Up to 31 devices can be addressed in a single system. Alternatively, if an unaddressed device is ready to send information, it can set the SRQ (Service Request) line low to alert the controller. The controller then polls each of the devices in turn (a serial poll) to determine which device has made the request. It is possible for more than one device to create SRQ at the same time, but the SRQ line remains low until all devices have been serviced. In effect, all the bus connected devices operate a "logic-OR" function on the SRQ line.

A very significant part of the HP-IB operation is the actual data transfer process itself. This is deliberately designed to be asynchronous so that system has great flexibility to handle devices with very different I/O speeds. The system adapts automatically to the rate of the slowest active device. The data transfer is controlled by a "three-wire handshake". The three wires, as shown in the diagram above, are DAV (Data Valid), NRFD (Not Ready For Data) and NDAC (Not Data Accepted).

In simple terms, the sender of a data byte (a controller or talker) waits until all the addressed (active) listeners indicate they are ready to receive data by setting NRFD line high (since the convention is "low-true", this means they are ready!). The line will not go high until all the listeners are ready (a "logic-AND"), so the process waits until the slowest device responds. The sender then places the data byte on the 8-bit data bus and, after a brief settling pause, sets the DAV line low to indicate the data is ready. When a listener has accepted the data, it lets its NDAC go high, however the bus NDAC line does not go high (data accepted) until all the listening devices have received the data and the slowest device sets NDAC high (another "logic-AND" process). NDAC stays high until the sender sets DAV high to conclude the data transfer process, and a new sequence can start for sending the next byte of data. No action in this series of events can be initiated before the previous step has been completed by all active devices, so data is transferred at the rate of the slowest addressed device.

The above is just a brief description of some key aspects of the HP-IB to help explain the South Queensferry products which extend its capability. There are a number of other control functions relating to talkers and controllers that go beyond what is necessary for this Chapter, however a useful summary of HP-IB operation is given in Appendix 8.2 in "*Telecommunications Measurements, Analysis and Instrumentation*"[3].

The HP-IB quickly established itself as an industry standard and by the mid-1970s was adopted by the IEEE in North America, issuing standard[4] 488.1-1975, "Digital Interface for Programmable Instrumentation". There was also the similar IEC 625-1 (1979). Other test equipment manufacturers used it too, though sometimes it was referred to as GPIB (General Purpose Interface Bus) to avoid using a major competitor's name.

The HP-IB was an effective interface with relatively high data throughput by the standards of 1970s and 80s. It was quoted as having data rates of 250 – 500 kb/s and optimistically up to 1 Mb/s for short connections. This made it attractive not just for instrumentation but also for connecting together computer equipment such as hard disc drives, tape drives, printers and terminals. Most HP equipment by the 1980s sported an HP-IB interface and the scale of its application expanded rapidly, in line with the growth of computer systems. Thus the market for HP-IB products became large, until the emergence of various serial high-speed LAN standards (such as Ethernet and Fiber Channel) finally displaced it.

---

[3] "*Telecommunications Measurements, Analysis and Instrumentation*" by Kamilo Feher/Engineers of HP, 1987, ISBN 0-13-902404-2, Appendix 8.2 pp. 304-313

[4] This standard defined the mechanical (connector) arrangement, the electrical details, and the all-important protocol layer which states how control operations and byte transfers are accomplished logically. This worked well, but it soon became apparent that incompatibilities were arising at higher protocol levels, especially because device designers structured their data to be transferred with differing coding and formatting – which often prevented one product from interpreting the meaning of another's information. To overcome this, a companion standard, IEEE 488.2 (1978) was later devised, known informally as the "codes and formats" standard.

Going Further

Once the HP-IB standard was established in the 1970s, the original 20 metre distance limit began to become a problem for some applications. One area was large-scale factory production and process control, and another was the monitoring large industrial plant such as power stations. Often, the measurement equipment would need to be installed close to the facility being monitored and the test results sent back to a central controller in an office environment. The distances involved could be several hundred metres, well beyond the HP-IB specification. Another application was telemetry or remote monitoring where the measurement equipment might be installed in a remote location, tens or hundreds of miles away from the control centre. A specific example, of particular interest to South Queensferry, was the performance monitoring of a distributed telecom network from a central network management control room. All of these applications required a means of HP-IB distance extension using a dedicated serial data link.

In the HP-IB system diagram shown above, imagine that the bus is divided in two between device B and C, so that devices C and D are remote from the main controller system, A. The 16-wire parallel HP-IB needs to be converted to a separate transmit and receive serial link and at the far end converted back to the parallel bus. However, the task is not as easy as it sounds because the original designers of the HP-IB never envisioned any form of serial operation. Key to this is the fact that the HP-IB interface has a transmit and receive function (using tri-state buffers which electrically can be in a “High”, “Low” or “Floating” state), with all devices on the Bus expected to communicate electrically through relatively fast metallic interconnection. This allows the logic-OR and logic-AND functions to be created in the Bus which are key to the handshaking and some control lines, as described earlier. Once a serial link of any kind is introduced, this direct logic functionality is broken and so has to be replicated in some way by the HP-IB extension process.

If high-speed serial data links were available, then it might be possible to interconnect the remote and local HP-IB systems in a fairly transparent manner, with only the inherent transmission delay (about 5 microseconds per kilometre) and parallel-to-serial latency causing delay in response times and handshaking.

In the 1970s, high-speed serial links were not usually available, and indeed any HP-IB extension using data modems over the analogue telephone network would have very slow operation at the data rates of the time, typically 1200 bits per second or less – nearly a thousand times slower than the fastest direct HP-IB operation. HP-IB extension over slow serial links, sometimes with bit errors present, was a considerable challenge to the designers of an HP-IB extender.

HP’s first product to offer this function was the 59403A HP-IB/Common Carrier Interface or CCI, introduced in 1975/76 by the Loveland Division in Colorado. Common Carrier simply refers to the US term for a telephone network, and the CCI was designed to operate over dial-up or leased-line modems at 1200 bits per second or less, with a CCI at each end of the link to interface to the HP-IB. Alternatively, the CCIs could be interconnected by two twisted-pair wire cables to provide bi-directional transmission up to 1000 metres, without data modems. When introduced, the 59403A sold for $1300, and there was also a companion box, the 59400A which converted serial RS-232 interfaces on display terminals and teletypewriters to HP-IB.

South Queensferry was an early user of the CCI in its first remote monitoring systems. In Chapter 5, I looked at the development of the ground-breaking 3745 Selective Level Measuring Set (SLMS) – the first microprocessor-controlled instrument in HP with fully automatic measurements, when introduced in 1975. It was also one of the first to have a comprehensive HP-IB interface. The Division was keen to take advantage of the automatic measurements in the SLMS to build a computer-controlled remote FDM monitoring system, as described in that chapter.

The first generation of HP-IB interfaces relied on a combination of TTL logic for the physical interface and handshaking, and usually processor firmware to handle messaging and addressing. Later, dedicated ICs became available to implement the interface logical functions, but they weren't around in the 1970s. Inevitably there were incompatibilities between the HP-IB implementation in different products, due to designers unwittingly taking "illegal" shortcuts in their designs, and this led to problems when building systems.

The 59403A CCI had some of these problems itself, particularly the handshaking and control sequences over a slow modem link. Robin Myles who worked on the HP-IB interface for the 3745 SLMS recalled that the CCI was on a constantly extending 26-week delivery from Loveland, partly because of all the problems South Queensferry was flagging-up as an early user.

There was also another fatal problem with the CCI in the remote HP-IB systems it was intended to service. The basic HP-IB standard (IEEE-488.1) didn't cater for transmission errors since none were expected on the short screened cables that it used for interconnection. But on analogue data modems, particularly on dial-up lines, bit errors were not uncommon, sometimes occurring in clusters due to transients on the telephone circuits. Although the CCI was designed to work with modems, it could only detect errors, not correct them. As a result, the errors could cause the remote CCI to hang-up, and the system would be disabled unless the remote unit could be manually reset. This was a major problem for a remote monitoring system that had to work 24 hours a day!

The Division needed a short-term fix if the remote monitoring project was going to progress. At the time, systems software was moving on to the HP 1000 computer, however the HP-IB card in the computer itself also had problems. It would hang-up too if there was a glitch. One of the South Queensferry systems engineers and software writers, Dave Warren, came up with a great piece of lateral thinking to write a special software driver, using the computer's 16-port multiplexer used for displays, and drive the remote CCI directly. This circumvented many of the CCI's problems, though not the lack of error correction. The Division was learning fast!

It was clear by now that the engineers at South Queensferry knew at least as much and probably more than anyone else in HP about how to operate the HP-IB over remote links. The Section Manager, David Dack, and the team realised there would be a good market for a product that provided robust HP-IB extension incorporating all the knowledge they had learned through the systems experience. Hence the first of a family of HP-IB extenders was born.

The 37201A HP-IB Extender

In April 1976, a new project team was set up, led by David Guest, to develop South Queensferry's first Bus Extender. The new product, the 37201A HP-IB Extender, would incorporate some key capabilities they believed would solve many of the problems of HP-IB extension that had been discovered during the previous two years:

- The three-wire handshaking would be optimised across the link to reduce the speed constraints of low-speed modem connections.
- The Bus Extender would incorporate error correction on the serial link so that modem errors would not affect HP-IB operation.
- Various timeouts and flow control would prevent complete system hang-up if there were a failure on the serial link.
- Point-to-multipoint systems would be possible when operating on suitably configured leased circuits, and an auto-dialler function in the Bus Extender would allow any arbitrary distant instrument site to be called at will on the public switched telephone network.

As explained earlier, the three-wire handshake is a key part of the HP-IB protocol, providing a reliable way of coping with the different I/O speeds of devices on the HP-IB Bus. Once the Bus is split by a serial link, the three handshake lines (DAV, NRFD, and NDAC) cannot meaningfully be sent to and from the far end because their logical interaction is lost due to the intervening serial link. However, the data byte and the five bus management lines, along with certain new control information, can be sent. The 37201A did this by assembling a two-byte unit comprising eight bits of management and "private" control information and the 8-bit data/address byte. Each Extender (one at each end of the link) would read the local HP-IB information into a serial buffer. Likewise in the receive direction it would read the serial data into a register and place it on the local HP-IB connector. The objective was to provide as near as possible a transparent link between the HP-IB subsystems at each end of the link so that the system programmer did not need to be aware of the serial communication between the Extenders.

As mentioned, the serial transmission used dual-byte (16-bit) data units assembled into variable length packets depending on the amount of data available. This was encapsulated with a header and other information including comprehensive parity check bits which formed the means for error detection across the serial link. When a packet was received error-free (as determined by the receive-end parity check), an acknowledgment packet was sent back to the transmit end and the next packet transmitted. If no acknowledgement was received, then the transmit end resent the packet until it was received correctly. Thus, if errors occurred on the modem link, the only effect would be a slightly slower operation.

In effect, the Bus Extenders had their own internal handshaking protocol on the serial link to adapt to varying HP-IB device speeds, error rate or link failure. It was, in fact, a central part of the product concept. For example, when two Bus Extenders were connected over the serial link, they would automatically exchange dummy packets to check continuity even without HP-IB traffic being present. The 37201A could then reliably flag a data link failure back to the controller at any time.

Another interesting application of the Bus Extender's serial link packet protocol was to replicate the effect of the three-wire handshake over the extended HP-IB. In regular HP-IB systems, the slowest instrument determines the speed of data transfer by holding-up the three-wire handshake as described earlier (the logic-AND process on the NDAC line). On the extended bus, this logic process does not work since the NDAC link and NRFD link are broken. Instead, flow control is regulated by the Extender's packet acknowledgement process, which operates between the two (now remote) separate handshake procedures. Thus, the handshake procedures at each end are separate for the most part and logically "legal", while still loosely coordinated by the data flow across the serial link. A slow instrument at the remote end causes the remote Bus Extender's buffer to fill and it then fails to return the acknowledgement packet to the local end. This is interpreted by the local Bus Extender which then uses the local NDAC line to hold back further data transfers.

The above examples give a brief insight into the many design ideas that went into the new 37201A HP-IB Extender. The objective was always to make the HP-IB extension as transparent as possible, so the system programmer could mostly ignore the presence of the Extenders, except for the potentially slower Bus operation.

Apart from acting as a transparent conduit in the HP-IB system most of the time, the Bus Extenders could also be addressed directly by the controller as HP-IB devices in their own right. This could be for obtaining status information on the overall extended system. Another application was to control an automatic dialler (attached to a separate connector on the Extender) so that a number of different modems at remote sites could be selectively called and then accessed via the local Bus Extender. Yet another feature was the ability to operate with modems connected to a Multipoint (or Multi-drop) leased telephone circuit. This Multi-drop arrangement was a useful way of connecting to several remote sites when switching between remote acquisition devices had to relatively rapid.

Another interesting idea was to address the local Bus Extender and program it into an "Idle" mode, which in effect disconnected the remote group of devices and the slow modem link connecting them. Under normal operation, the overall system would operate quite slowly because of the handshaking across the remote link (all messages need to be transmitted to the remote site), even when transactions were only between the locally connected devices. The "Idle" mode allowed the local Bus to run much more quickly, as high as 2400 bytes/second compared 38 data bytes/second when operating over a 1200 b/s modem link.

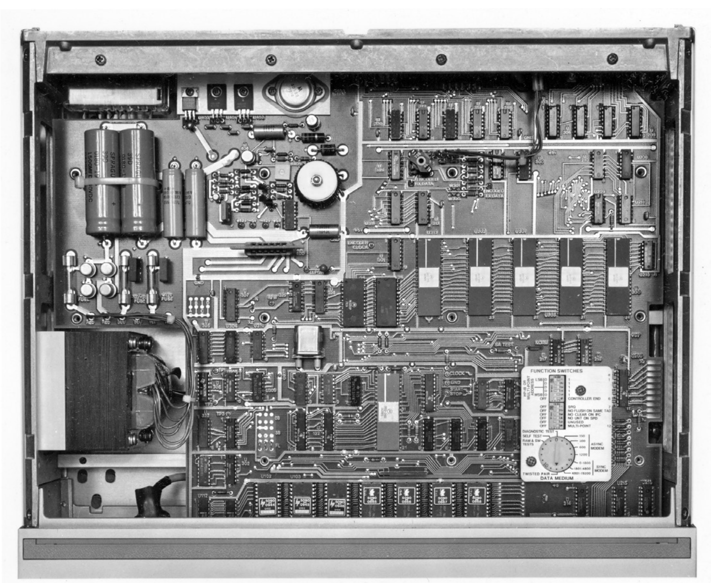

As with other HP-IB products of the time, the 37201A Bus Extender used a combination of discrete TTL logic ICs to implement the physical Bus and a processor (in this case a Motorola 6800) to handle the more complex messaging, flow control, error control, status and handshaking. Most of the higher level functions in the Bus Extender were implemented in firmware which allowed updates and enhancements.

A more complete description of the design philosophy of the Bus Extender was published in the August 1979 edition of the *HP Journal*[5].

David Guest, the project leader, recalled how he and the small design team were left alone in a corner of the R&D lab to start from cold, learn and just get on with it:

"*Those were the halcyon days when 'HP Invent' meant something, and engineers didn't have to account for every minute of their day on a timescale chart created by someone else!*"

David also remembered there was some debate about the implementation of the final product:

"*The first mock-up of the 37201A comprised several plug-in boards on a motherboard, as then commonly adopted at Queensferry. Suddenly, we realised this was unnecessarily complex and expensive, and we decided to place everything on a single larger board. The product support manager couldn't see the sense of this at all. How could it be serviced if boards couldn't be exchanged? A minor turf battle ensued, but we stood our ground on the basis of much increased reliability and lower manufacturing cost. In reality, simplicity won out, and very few failed in the field. The single board was accommodated in a rack-wide cabinet, in recognition of the HP-IB Extender as a system component.*"

Like the earlier CCIs, the 37201A Bus Extenders could also operate the serial link over two pairs of twisted wires to a distance of 1000 metres. The transmission rate was 20 kb/s compared to 1200 b/s maximum rate over the data modems, raising the data rate from around 38 bytes/second to 775 bytes/second. Line coding and balanced transformer coupling improved the immunity to interference on the twisted pair transmission. This diagram, from the HP catalogue in the 1980s, shows the 37201A in an HP-IB system with the two options for serial transmission:

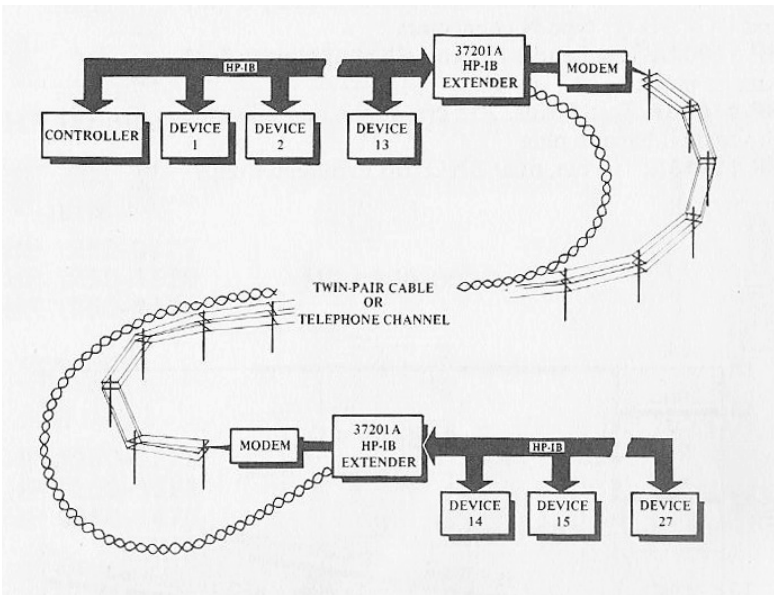

Point-to-point connection using twin twisted pair cable or full duplex modem link.

[5] "An HP-IB Extender for Distributed Instrument Systems" by David Guest, HP Journal August 1979 pp 26-32. http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/1979-08.pdf

The 37201A Bus Extender was introduced in early 1979 at a US list price of $1840. It sold at a fairly steady rate of 350 to 400 units per year to all parts of the world, with perhaps 60% going to North America. Apart from some minor enhancements and modifications it remained largely unchanged throughout its production life, confirming the original recipe and that the ideas of the South Queensferry engineers did indeed stand the test of time, unlike some of the other products that attempted this task! It was finally discontinued in 1991, by which time the price had risen to $2800. It had total sales of 4200 units amounting to around $10M revenue.

The 37203A HP-IB Extender

Recalling the genesis of the second Bus Extender product from South Queesnferry, Simon Murray, one of the development engineers commented:

"*With the 37201A, the local and remote HP-IBs ran somewhat independently, especially in that the three-wire handshake operated separately at each end with little end-to-end coordination. As a result, a few corner-cases of the HP-IB protocol wouldn't work adequately and had to be circumvented by adjustments to the controller programme. But it had to be that way to get half-decent throughput over the data modems of the day (1200 b/s was bleeding edge back then). Full end-to-end coordination would have required many more transmissions of data across the serial link, resulting in quite unacceptably slow system operation. Every time we hit a corner case, we would say 'if only we weren't limited by modem speed, we could do this so much better'. So the following 37203A was really all those 'if onlys' turned into hardware – wishful thinking made real!*"

The reality was that most applications for HP-IB extension did not need the great distances which were possible with modems and telephone links. Moreover, much faster operational speed was required for most system applications, particularly involving computer storage and peripherals.

In those days, HP engineers were allowed some freedom, and under-the-bench projects were often tolerated. The 37203A was devised unofficially during the latter stages of the 37201A. A quick breadboard version worked in most respects, and was then exposed to the light of day. David Guest stopped Bob Coackley (R&D Manager) in the corridor one day to reveal that they now had a fast solution to HP-IB extension. When he suggested this should be a regular new product, Bob's immediate reply was: "*Well I suppose you had better do it then!*" No great long market investigations and interminable meetings, just "Do it".

Perhaps it was just as well the 37203A remained "under the radar", as another team over at Data Systems Division (DSD) in Cupertino, California, had got there first. In the early 1970s, DSD[6] was famous for HP's leading computer products such as the HP 1000 and HP

[6] DSD was part of the original HP Data Products Group that produced the early 2114 and 2116 computers. In the mid-1970s it spawned the General Systems Division that took over the HP 3000, and the Boise

3000 systems as well as the mass storage devices and peripherals. They needed a high-speed Bus Extender for their computer systems and developed one in the late 1970s, pretty much at the same time as Queensferry was developing the 37201A.

The DSD product, the 12050A, was a high-tech creation using HP's new fibre optic LED[7] transmitter and receiver (developed by Optical Components Division (OCD)), along with multimode fibre cable for the serial link. The data rate was around 1 Mb/s on the fibre, which allowed the HP-IB to operate at up to 20 kilobytes/second, 25 times faster than the 37201A, although the maximum distance for the 12050A was limited to 100 metres compared to the 1000 metres on the 37201A with twisted pairs.

The electronics was also high-tech as it used some proprietary chips. DSD had been working with HP Labs on a new family of Silicon on Sapphire CMOS processor chips for use in their latest computers. The designers used a 16-bit processor chip and also the PHI chip (Processor to HP-IB) from this family to accomplish the high-speed extension. The price was also high-end at $1950 per unit and the product was introduced in August 1979, much the same time as Queensferry introduced the 37201A. More information on the design of the 12050A was published in the *HP Journal* of December 1979[8].

Undeterred, the South Queensferry team got on with their "official" under-the-bench project. It was the opposite of the high-tech complexity in DSD's 12050A, as it was based on the philosophy that the simpler it was, the faster it would go. The central secret of the new 37203A was to stretch the HP-IB three-wire handshake state machine over distance so that the successive states of the handshake were closely coordinated and truly the same everywhere. This became possible because the serial link was now on a coaxial cable and capable of fast operation.

It was an ingenious design, as Simon Murray recalled:

"*No processors, just logic. And as ever, although the idea was superbly simple, its execution was fiendishly clever in correctly coordinating all the 'legal' state transitions of the HP-IB standard, not just the most obvious handshake procedures.*"

A US patent for this design was filed by David Guest[9] and Peter Roubaud in May 1983, Patent 4,451,886 "*Bus Extender Circuit for Data Transmission*".

The designers selected coaxial cable for the serial link as it was economical, easily sourced and installed, and provided good transmission bandwidth. A single cable was used for both directions of transmission (half-duplex), further simplifying installation. As with the earlier extender, the transmission bytes were assembled in a frame and sent with modest error-checking overhead (in this case Cyclic Redundancy Checksum). When a frame was received error-free, and acknowledgement was sent back, otherwise the send end would assume an error and retransmit the frame. Of course, errors were very infrequent

---

Division in Idaho, famous for the 7900 series disc drives and later the HP laser printers. For more on HP computers and divisions, see http://www.hpmuseum.net/index.php

[7] Light Emitting Diode. The first fibre optic systems in the 1970s used LEDs (both infra-red and visible) as the optical transmitter, before the days of solid-state laser transmitters

[8] "High-Speed Fiber Optic Link Provides Reliable Real-Time HP-IB Extension" by Robert Grady, HP Journal December 1979, pp3-9. http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/1979-12.pdf

[9] David Guest filed a second patent in 1987 (US Patent 4,852,043) which covered an invention for dealing with remote HP-IB failure (a truncated Bus), which was used in the later 37204A HP-IB Extender

compared to the modem links used by the 37201A. The coax link turned out to be "bullet-proof" in the face of interfering electromagnetic fields – careful attention to hardware design ensured it worked flawlessly even when wrapped three times round a turbo-generator in an electricity power station!

The 37203 was also available optionally with a fibre-optic serial link, similar to that used in the 12050A. Some customers thought they needed optical communication to avoid electromagnetic interference, but they probably didn't. Simon Murray commented:

"*We sold more optics than expected so it was kept going as an even more profitable option. It was partly a "Faustian" deal to keep the Americans off our patch, since OCD were desperate for sales of their optical Tx/Rx modules.*"

In due course, OCD upgraded the range of their components from 100 metres to 1000 metres, and this was fitted to the 37203. David Guest recalls an amusing incident during the development, when Bill Hewlett made an impromptu visit with the section manager, David Dack. Bill's first comment was "*There seem to be a lot of Davids round here*", to which David Guest's response (alluding to Bill and Dave) was, "*I thought it was a good name to have in this company!*" "*Yes, should go far*", was the response from Bill Hewlett, who it seems knew more about the new LED transmitter from OCD than we did.

The new optics did provide a speed advantage at distances over 500 metres, because of the high-frequency losses in the coaxial cable, but for most applications the coax gave comparable performance as this table from the 1985 HP Catalog shows:

Distance	**Max HP-IB transfer rate (kbytes/sec)**	
	Coax Link	**Optical Link**
Short	50	50
250m	40	39
500m	14.2	32
1000m	2.75	25

The world's first properly transparent HP-IB extender was introduced in July 1980, in two versions. The standard product in an HP half-width cabinet was the 37203A costing $1295, and a card version for the HP 1000 computer, designated 37203L, sold for $1030. The optical link option cost an additional $470. Prices did not change significantly during the life of the product.

Compared to DSD's 12050A, the 37203A was twice as fast and a lot less expensive. It also seemed to work better judging by a story recounted by Simon Murray:

"*Derek Milne, the product support engineer, took an early 37203A pair to the USA. A field sales guy took him to see a disgruntled customer who had bought the DSD boxes but couldn't get them working. They were lying to one side with their fibre cables still attached. Derek swapped the cables over on to our boxes, and hey presto, it worked perfectly. Cue, one amazed customer.*"

The 12050A rapidly faded from view and disappeared completely in the 1982 HP Catalogue.

The market for a short-range high-speed Bus Extender was huge, and the 37203A unit sales ramped-up rapidly to ten times the level of the 37201A. A key market within HP was the remote connection of printers to the HP 3000 computer.

Production volumes of 4000 units a year and more were unheard of at South Queensferry before then. The 20,000th unit was shipped on the 19th September 1986. The 37203A was in production until 1988 when it was superseded by the 37204A, described next. Here it is pictured in a promotional advert in the 1980s. In total, around 24,000 units were produced with revenues of $34M (assuming 25% optical option) – not bad for an under-the-bench project!

The 37204A Bus Extender (NMS T.1986.617)

The 37203A had been an enormous success with high production throughput. Attention then turned to what could be improved. Technically very little needed to change as the 37203 seemed to do the job perfectly, so the focus was on reducing production cost and assembly time, and enhancing reliability. The key contribution in the new product was to condense a whole board of discrete logic ICs into a single custom "gate array". For this task the team turned to another R&D engineer, John McElroy, who had experience of chip design. John had previously designed South Queensferry's digital filter chip for the 3776 PCM Test Set (see Chapter 8).

The gate array had 1600 gates, then a considerable number on a single chip. Committing to a gate array was a big step especially as simulation was not as good as it would become later. To avoid expensive reengineering, a discrete-logic breadboard was wired up first before committing to silicon. John McElroy described the design:

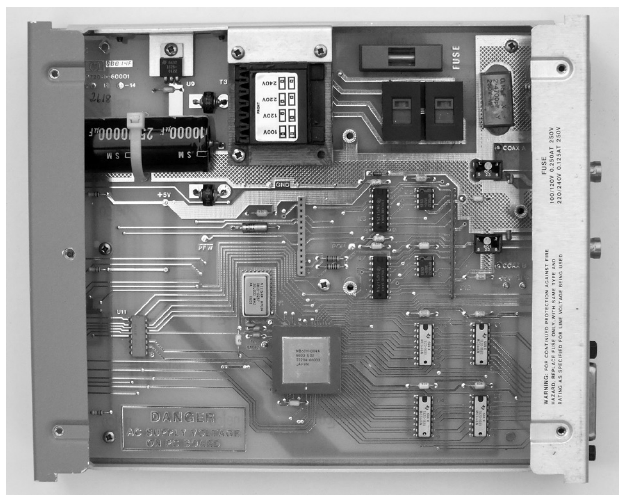

"*David Guest focussed on the HP-IB state machines and I worked on the serial communications logic. We built a full scale model using programmable logic arrays and I then developed the Fujitsu gate-array (visible lower centre in this photo). The final product had one gate-array, some HP-IB driver chips and two coax line interface circuits.*"

David Guest recalled that, just before the 37204A was launched, a curator from the National Museum of Scotland was visiting the factory looking for technical artefacts to acquire. He was very interested in the concept of the gate array and its development and application. As a result, the 37204 along with its gate array and breadboard simulator had the rare distinction of finding its way into the Museum collection before it even shipped! These items have NMS Accession Numbers T.1986.614 - 617.

The 37204A also had some further advantages over the 37203A. The range was increased to 1250 metres on coaxial cable and eventually up to 3000 metres on optical fibre. It was also capable of multi-point operation whereby Extenders could be daisy-chained together so each remote site only needed one 37204. It had a patented Chain Truncation facility which allowed the computer site to continue operation in the event of remote power failure or link failure. Daisy-chaining did reduce the byte transfer rate compared to the point-to-point maximum of 60 kbyte/second.

Another improvement in the new product was the packaging. Whereas the 37203 had used standard HP cabinet hardware requiring some assembly work, the 37204 used a novel low-cost plastic housing devised by product designer, Tony Cowlin. Sometimes referred to as a "clam shell", it clipped together round the printed circuit board in seconds. The production time for the complete unit was reduced to a few minutes, and the assembly worker usually tested its operation as well. This was the first time one person handled the production process from beginning to end.

The 37204A was a direct replacement for the 37203A and was introduced in February 1986, taking over from the 37203A completely by the end of 1988. As Mike Kerr recalled, the timing was perfect as a Japanese competitor had just introduced a competitor for the 37203 hoping to get a share of the market, but the new 37204A killed it dead. Again volumes were initially around 4000 units per year. At introduction it was priced at $925 (a reduction on the 37203 which sold for $1285) but the profit contribution was higher. The price was gradually ramped-up over the following years when no competition appeared, reaching $1.5k in 1993 and $2k in 1997. The Division was certainly "milking" the market as the volume tailed off. However, only in the last year or two did the 37204A match the price of DSD's 12050A from nearly 20 years earlier!

The 37204 was in production until 1999. Volumes began to drop in the 1990s as HP-IB was gradually superseded by other control interfaces such as Ethernet, particularly on equipment such as printers. By 1997, the volume had dropped to just over 1000 units, compared to 4000+ units per year a decade earlier. A total of around 31,000 units were shipped bringing in revenue of $46M.

One final idea that never saw the light of day was to put the 37204A gate-array chip in a coax-only stretched HP-IB connector with a wall-plug power supply module. A mock-up was shown at a Corporate Review but it was never progressed, mainly because there was little market incentive. The 37204A already made very good profits and there was no real competition.

Extender Retrospective

For nearly 20 years, South Queensferry completely dominated the world market for HP-IB Extender products. Some competitors occasionally came on the scene, but they had little impact as the products weren't as well thought through as the HP Extenders. It wouldn't be unrealistic to say the Division had 80 to 90% of the market, and of course the HP branding was always a significant advantage since the Company originated the HP-IB and was a major supplier of HP-IB controlled products in its instruments as well as its computer portfolio.

Between 1979 and 1999, South Queensferry shipped over 58,000 units with total revenue of over $90M. Perhaps more significant, this was very profitable revenue and it is likely that the gross profit margin was 30 to 40%, giving a net income of well over $30M. The Bus Extender family had the highest return on investment of any of the South Queensferry product developments.

In the late 1980s, around a third of the total Division profit was coming from the Bus Extender family; an astonishing thought considering the products were invented by a group of three or four engineers, with a similar number of staff making the units in production.

The success of the Bus Extenders was if anything an embarrassment to the Division Management, whose strategy was, of course, based on growing a profitable line of Telecom Test products. The Extenders were supposed to be a sideline, but were making an awkwardly large proportion of the Division's profits, so attention couldn't be drawn to this embarrassing fact. Mike Kerr (Product Support Engineer) enjoyed drawing an analogy with brothel-keeping – management were happy to take the profits, but kept quiet about the source!

The unit sales profile for the three main Bus Extender products is shown in the graph below. Given the high market share, it is probably also a reasonable reflection of the growth and decline of the HP-IB market as a whole.

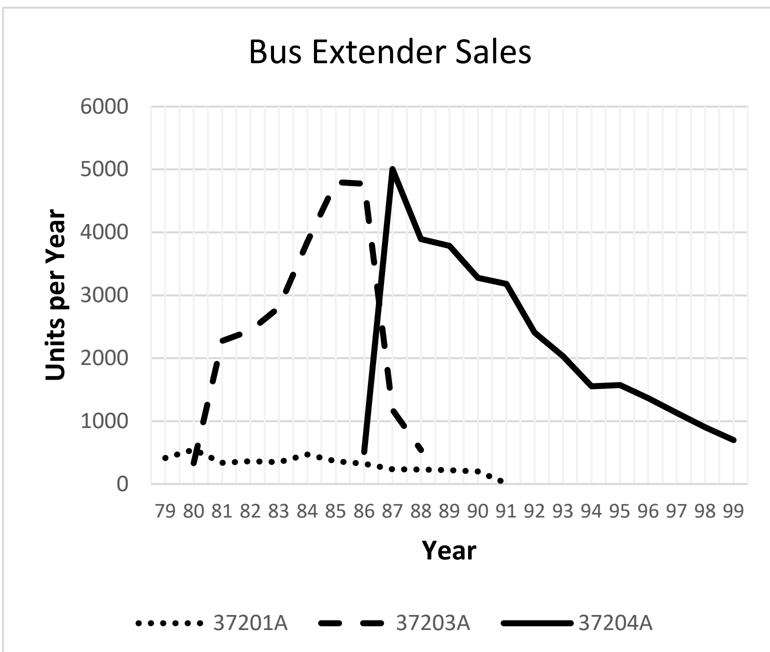

South Queensferry's Bus Extenders were a perfect example of the power of ideas to generate wealth. In hardware terms there was nothing exceptional in the Extenders (apart from the gate array) and the components were all readily available industry-standard items. The value of these products was entirely in the ideas and the logical thought processes embedded in this hardware.

This intellectual capital was never really challenged either inside or outside HP.

The Modems

The second part of this chapter focuses on another group of products somewhat related to the Bus Extenders and again directed at a much wider market than Telecom Test. These were the voice-data modems.

From the 1960s to the 1990s, they were the main method for transmitting data long distance over the telephone network in the days before direct digital services became available. Whether it was a computer network or just a fax machine, everyone needed data modems. The principle was that the digital data would be modulated onto a carrier to produce a band-limited signal that would pass through the 3.1 kHz telephone voice channel. It was demodulated at the far end to regenerate the original data stream, hence the name modem – short for modulator/demodulator. The more sophisticated modems used more complex modulation and could transmit data at a higher rate, but were also more vulnerable to imperfections in the telephone circuit, so required adaptive equalizers. See Chapter 6 on Telephone Line Analyzers (pages 105 to 108) for more information on modems and the effects of telephone line impairments.

For South Queensferry, the Bus Extenders and the Modems had something in common: both started as a response to an internal need. In the case of the Bus Extenders, the Division needed to find a reliable solution for remote control so it could develop its FDM surveillance systems. In the case of the modems, the HP computer group was having problems supplying integrated computer networks which relied on third-party modems for the communication links. They turned to South Queensferry, with its knowledge of communications networks, to come up with an HP solution.

Robin Myles, who around 1980 was handling product marketing for systems, described how this happened:

"*It was at a Division Review at which Ed McCracken, the GM of General Systems Division who made the HP 3000, was present. I think they were having a lot of hassle with customers who were using modems to connect terminals to mainframes and also computer-to-computer. Because these modems were provided by a third-party, whenever there was a problem there was some finger-pointing. We were regarded as a good product division to buy in a modem and badge it as HP, so the computer sales force could offer a single-vendor solution.*"

The Paradyne Deal

In late 1977, South Queensferry negotiated an OEM[10] contract with Paradyne of Largo, Florida to supply modem boards which would be built into a standard HP system II cabinet with HP branding, literature and post-sales support. In the 1970s HP itself already used Paradyne modems operating at 2.4 and 4.8 kb/s to link up its international computer network called COMSYS. Long before the age of the Internet, COMSYS provided an internal electronic messaging system between all the sales offices and divisions as well as a means for transmitting orders and other financial information.

[10] Original Equipment Manufacture is the term used in the industry to describe component manufactured by one company and then rebranded and sold by another company under their name, sometimes also described as a Value-Added Reseller.

The Division produced two OEM synchronous modem products, the 37210T which operated at 4.8 kb/s and the 37220T which went to 9.6 kb/s. The Paradyne modem cards were fairly high performance with adaptive equalizers to compensate for imperfections in the telephone line. The 37210T (which conformed to the CCITT V.27 standard) could operate over leased lines or dial-up circuits[11] and could reduce its rate from 4.8 kb/s to 2.4 kb/s (8PSK to 4PSK modulation) if line quality was poor. The 37220T was for operation over leased lines only at 9.6 kb/s with a fall-back rate of 4.8 kb/s and used 16QAM modulation (CCITT V.29 standard).

While the product development was simple, the product marketing was not. Before deregulation of the telecom industry in the mid-1980s, most national telephone companies were state-owned monopolies. Each country had its own regulations on what could or could not be connected to the telephone network. Some operators considered a data modem as part of the operational equipment and the customer had to use modems supplied exclusively by them. For others, end-user equipment could be connected but had to be registered with the operator for connection approval, sometimes with specific restrictions on the line interface or bit rate. A further complication was that some authorities allowed connection to leased lines but not dial-up circuits.

It was a minefield, but the South Queensferry marketing team had to get approval anywhere the new modems were to be sold. Without it, HP didn't have the single-vendor solution. In his memoirs, Hank Taylor[12], the architect of HP's COMSYS systems, highlighted the difficulties of setting up modem links internationally:

"*Outside the US, the governments of each nation kept control over all telephone equipment and transmission lines. Some countries were more cooperative than others. As we established dial-up connections to the tougher locations, their restrictions were painful, allowing only 1200 bps modems, which were manufactured under their direct control. Geneva, Switzerland was a primary hub for all of HP Europe's transmissions and was a very high volume location, but the Swiss PTT was very restrictive. We negotiated with them for some time to install a faster non-Swiss modem and their answer was always, 'No!' Finally, in exasperation I told them that unless we could install a faster modem, we would move our European Headquarters to another country that was more cooperative. They said, 'Give us a minute.' They came back and said that we could install a faster modem, but that it would have to be called a test installation and they wanted it located underneath the raised computer floor so as not to be visible.*"

The 21st century reader will be astonished at the difficulty of setting up a computer network in those days. They will probably be even more astonished that any useful information could be exchanged across a 9.6 kb/s link, let alone 1.2 or 2.4 kb/s which was more typical! However, one has to remember that in 1980 even quite large computers had a main memory of only a few hundred kilobytes, and computer messages were usually short strings of ASCII alphanumeric characters.

The new modems, which were known in the factory as the "Paradyne Modems", were introduced in 1980 with list prices of $4.3k for the 37210T and $6.1k for the 37220T, plus additional costs for various options which were obligatory in some networks.

---

[11] Usually leased lines were 4-wire circuits allowing full-duplex operation with 4.8 kb/s in both directions simultaneously, whereas dial-up was 2-wire with half-duplex operation – 4.8kb/s in alternate directions.

[12] To read his complete memoirs, see http://www.hpmemory.org/timeline/hank_taylor/experiences_01.htm

In the October 1980 Computer News, HP's computer sales-force newsletter, the new products were listed as approved in the USA, Canada and several European countries.

The Paradyne Modems (shown here under test) didn't achieve high volumes, probably 10 to 12 units per month, mostly sold as an adjunct to HP 3000 computer sales. They were high-end, and as Gordon Reid, product manager for the modems, recalled:

"*It was difficult for an instrument division to sell via the HP computer sales force because normally we had little contact with them. We didn't get much access to train staff and motivate them to sell our products, and I think their lack of understanding meant they were a bit wary and didn't try very hard to push the 37210/220T.*"

It has also been suggested there were problems on the supply side, with erratic delivery of modem boards from Paradyne and technical faults. For whatever reason, the modems were discontinued by early 1983 and not replaced, so presumably the HP Computer Group had found an alternative solution for systems sales. Total sales were 281 units for the 37210T and 367 for the 37220T.

Modems Made in Scotland

It is clear from South Queensferry's 1980 Strategic Business Plan that the Division had ambitions in the modem market. The Plan forecast orders from Modems and Bus Extenders rising from $2.28M in 1980 to $14.35M in 1984, a significant part of the total business. A year earlier, the 1979 plan seemed even more ambitious, with modem revenue rising from $200k in 1979 to $7.7M in 1982! Over the years, it has to be said, these South Queensferry Intermediate Range Plans (IRPs) became notorious for being completely "pie in the sky", however for modems it did show the intent.

Worldwide, it was a very large market in the early 1980s as modems were more or less the only way to build Wide Area Networks (WANs) for computers, before direct digital services became available in the late 1980s. The annual global market was probably worth at least $300M[13]. South Queensferry wanted some of that.

I remember Jim Rigby, who was the Financial Controller at the time, getting fired-up about the opportunity, no doubt encouraged by the rapidly rising sales and excellent return on investment for Bus Extenders, which were seen as a similar market. He advocated that the site should set up a new division to manufacture these products separately from the instrument business which was based on low-volume batch production. He probably imagined himself as Division Manager.

[13] One of the major modem suppliers, Racal-Milgo, reported revenues of $100M in 1979.

It was a reasonable proposition since the cost accounting model would need to be different as would the production process for high-volume modem manufacture, which would require investment in automation and Just-in-Time (JIT) stock control. To be competitive, it would also require considerable capital investment in custom large-scale integrated circuits.

The R&D design team certainly had the expertise to develop advanced modems. There was considerable experience in logic design and processor software, and the Division was very familiar with telephone circuit parameters. There was also expertise in advanced digital signal processing in PCM test and with the development of the digital filter chip in the early 1980s (see Chapter 8 on PCM Testers). In fact in the early 1980s the lab developed the 4948A In-Service TIMS which was in effect a glorified high-end modem "with knobs on". (See Chapter 6 page 105)

On the downside, although the market was large, it was also very tight. There were the regulatory restrictions mentioned earlier and it was pretty competitive. Several large manufacturers, including Paradyne, Codex and Racal-Milgo, had been in the market for 15 to 20 years and were entrenched. They had also invented several landmark products raising modem bit rates and improving digital adaptive equalizers. How was South Queensferry going to attack this market, rather late in the day? And how would they get the products to market through the HP instrument sales force?

Nothing came of any of this. Perhaps reality dawned that it was going to be difficult to make money out of this business. It seemed like a similar market to the HP-IB Extenders, however with the Extenders, South Queensferry had got there first and always had the inside track with HP's large captive market. Although there were further technical advances in modem design in the 1980s with more sophisticated modulation and coding schemes and more powerful adaptive equalizers, there must have been those at South Queensferry immersed in the new digital network revolution who would have believed that time was running out for the voice-data modem market. Nobody would then have foreseen the enormous boom in dial-up modems in the 1990s to serve the new Internet age. That was still more than a decade in the future.

However, the team didn't give up on the modem market. Instead, they seemed to return to the original business proposition of providing modem networking equipment for HP's growing computer business as well as the Division's remote testing systems such as RATES.

37212A/B Dial-Up Modem

Rather than focusing on bit-rate heroics to achieve the maximum throughput on a point-to-point leased telephone line, the new product was aimed at the dial-up market using the 2-wire circuits of the switched telephone network. The 37212A was a 1200 b/s full-duplex modem, meaning it could send 1200 b/s simultaneously in both directions. It conformed to the Bell 212A and CCITT V.22 standards which used a version of 4-PSK modulation on two frequency-separated carriers within the telephone channel to minimize interference between the go and return paths.

The 37212A positioned itself as an intelligent modem with processor control, similar to the contemporary Hayes "*Smartmodem*", and there was a secondary data channel for controlling the modem. It had auto-dialing (using pulse or tone) and auto-answer and could configure automatically to the modem at the far end. Thus it could be used to build point-to-multipoint data networks using the switched telephone network. The non-volatile internal memory could be used for storing telephone numbers and log-on sequences for faster operation.

There was also a "Systems Modem" version with the 37212A on a plug-in card (37213A at $1100) and a card-cage (37214A at $1700) and some other interface cards. It sold quite well, with 1670 card cages and 2800 modem cards (37213A). It began shipping in March 1983 and continued until 1990/91, mainly because it was part of the RATES system.

The 37212A itself was introduced in November 1983 and sold for $1050. It was in production until February 1988 when it was replaced by the 37212B. A total of 4950 units shipped, and of these 1750 were sold in the UK, many probably associated with the BT RATES contracts.

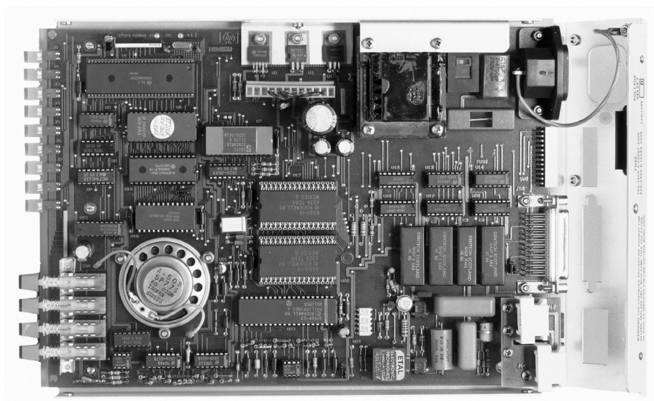

The replacement 37212B (shown here along with its internal circuit board) was a significant improvement on its predecessor. It used the CCITT V.22bis standard introduced in 1984. This was similar to V.22 with the dual carrier full-duplex operation but instead of using 4PSK it used 16QAM modulation, which doubled the bit rate to 2400 b/s in the same bandwidth[14], along with an adaptive equalizer to correct line impairments. The new modem was backward compatible and would automatically switch to earlier standards at 1200 b/s, and even 300 b/s.

Another significant improvement in the new modem was built-in error correction. It used the Microcom Networking Protocol (MNP) which assembled the incoming data into packets and sent them with a checksum for error-checking. An acknowledgement message was interleaved with the data in the return direction to confirm the packet had been received error-free.

The 37212B was a popular product, and selling at $920 (US list) it was cheaper than the A model but offered more performance. It started shipping in October 1987, with a forecast of 250 to 300 units/month, and was in production until April 1991. A total of 9000 units were sold, 80% in the USA, bringing in over $8M.

[14] Whereas the V.22 modem encoded two bits on each symbol resulting in a 600 baud symbol rate at 1200 b/s, the 16QAM modulation encoded four bits on each phase state or symbol and so could transmit twice as much information at the same baud rate.

Overall, the various products in the 37212 modem family produced revenue of about $20M between 1983 and 1991. It was a tiny share of the worldwide modem market but the products made a useful contribution to HP's computer sales activity and particularly to South Queensferry's RATES business, described in Chapter 7.

Postscript – the 37230A Short Haul Modem

To end this chapter, here's the story of a product that was a bit of an orphan. Most modems, including the ones discussed so far, were designed to operate over multiplex carrier sections in the trunk network (FDM analogue or PCM digital) which limited the bandwidth to 3.1 kHz. Hence all the complex phase modulation schemes which were developed to get the maximum bit rate through the restricted bandwidth. If the circuit was entirely on a pair of copper wires, for example the local loop from the subscriber to the exchange, then potentially much more bandwidth could be available[15].

The 37230A Short Haul Modem was designed to operate over these metallic circuits, either leased from the telephone company or installed privately. There was no need for any complicated modulation, so this type of modem was also sometimes referred to as a Baseband Modem. It converted the incoming data received on an RS232C interface to a line-coded digital stream adapted to the characteristics of the metallic circuit and eliminating DC components. The 37230A had an automatic equalizer to compensate for variations in the telephone line characteristics. The distance it could work depended on the transmission rate and the type of circuit: over 20 miles at 2.4 kb/s to 9 miles at 19.2 kb/s.

The 37230A was introduced in 1981 and cost around $1200. It made modest sales of 10 to 20 a month and had a short life, being withdrawn in 1985 with total production of 760 units. The problem was that nobody in HP at that time sold stand-alone communications equipment. The instrument sales-force, with whom South Queensferry normally worked, didn't deal with the right customers. The computer sales force probably did, however they lacked the technical knowledge to sell it. In fact calling it a "modem" may well have confused them, so they steered clear of it in case they sold the wrong thing. At least they were on safer ground selling the 37212A/B modem family as it would work anywhere. Sadly, the Short Haul Modem didn't really fit in to HP's portfolio. Perhaps the orphan might have done better described as the 37230A RS232 Extender and sold alongside the 37203A HP-IB Extender.

Acknowledgements

The following former South Queensferry employees have contributed to this chapter: David Guest, Simon Murray, Robin Myles, John McElroy, Mike Kerr and Gordon Reid.

[15] This was taken further in the 1990s with the V.90 56 kb/s dial-up modem, and later with broadband access using ADSL

13

Chapter Thirteen

Dynamic Signal Analysis

This story takes us back to the early days of HP at South Queensferry in the late 1960s. When the new Division was set up in 1965, most of the activity (as it had been at HP's first UK operation in Bedford), was focussed on transferred products which were designed at other HP divisions and manufactured under licence in the UK for the European market. But right from the beginning, HP intended the new Division to have its own design team developing a unique range of products for the world market – a worldwide product line. The question was what should that be?

In Chapter 2, we saw how Peter Carmichael and Finlay Mackenzie came from Ferranti in 1965 with the idea for a Microwave Link Analyzer. At the same time another team from Edinburgh University led by Gordon Roberts came to HP with ideas for a new family of instruments for testing mechanical and control systems using random noise.

After graduating in 1954, Gordon Roberts did a PhD at Manchester University on noise in non-linear systems. After a few years in industry, he returned to academic life at Edinburgh University in the early 1960s, lecturing on control theory and investigating the use of noise signals in systems evaluation. His colleague, Brian Finnie, had come to Edinburgh University in 1962 to do a PhD on real-time correlation with noise signals, working in the research team headed by Gordon. This expertise formed the basis for the proposal Gordon and Brian took to HP in 1965.

This interest in noise test signals and analysis requires some explanation. The opening paragraphs in the September 1967 HP Journal sets the scene rather nicely:

"*Almost every natural and man-made system is subject to random disturbances under normal operating conditions. Consequently, it is often appropriate and sometimes essential to test a system with random test signals rather than the sine waves that are so familiar to electrical engineers. Many of the areas of application for random test signals* *lie outside the field of electrical engineering. Examples are biomedical phenomena, vibration, structural analysis, aerodynamics and seismology.*"

As Gordon Roberts pointed out, injecting large sinusoidal test signals, or disturbances, in a working system was impractical:

"*The major problems in automatic control systems, such as nuclear reactors or steering a 250 ft radio telescope dish, were that you could not safely inject sufficiently large sinusoidal disturbances to accurately measure the system's frequency response. The systems were noisy, and the test signal responses were hidden in the background. The possibility of using random noise as a test signal was attractive. I had a contract at Edinburgh University from UKAEA*[1] *to develop a test method for the Dounreay Fast Breeder Reactor in Caithness, and Brian Finnie worked on a random noise solution for his PhD.*"

Gordon gave some further historical perspectives:

"*As a result of wartime development in armament technology, the specifications were getting tighter every year, especially in the new fields of guided missiles. The demands on the design engineer for better target-tracking, can be translated into "better frequency response" of the missile's aerobatic performance. The classic approach was to assume linearity, and just get the control system poles and zeros in the right place. In real life, gross nonlinearity was the norm, signals were hardly ever sinusoidal and there was no established theory as to how to design the optimum system.*

"*So in the mid-1950s there was much research into better methods of attacking the problem, using random signals which more closely matched the real life situation. To begin with, random signals with known spectral and amplitude probability values were not easy to generate using analogue techniques. The target specification was a true Gaussian probability distribution, and a spectrum going all the way down to zero frequency. In parallel with this was a quest for analytical methods better suited to the random signals and the nonlinear systems.*"

Apart from being qualitatively more appropriate, what are the other advantages of using random noise as a test signal?

Many of the large control systems or physical structures being tested have a relatively slow response time, or in electrical equivalent, very low bandwidth. This means if sinusoidal test signals were used, they could be well below 1 Hz, even milli-Hertz (one thousandth of a Hertz). It takes a very long time to measure just one cycle of the test signal, and many cycles may be required if noise is present.

To measure the frequency response of the system – output versus different test frequencies, would take an inordinately long time, during which the characteristics of the system itself might change. Random noise (or white noise) is effectively a spread spectrum signal with all test frequencies present simultaneously, so the measurements can in theory be done in parallel, greatly reducing measurement time.

[1] United Kingdom Atomic Energy Authority

Correlation

None of this immediately solves the first problem that a safe test signal could easily be buried in the background system noise, noise which of course is unrelated to the random noise test signal. These two noise signals are said to be "uncorrelated". The background noise is truly random and unpredictable, while the desired measurements are being made on the known noise test signal after it has passed through the system under test, so the input and output are correlated. If the measurements are made over a period of time, the correlated test signals will yield a concrete result (the system response), whereas the uncorrelated background system noise will eventually average to zero.

A way of testing for correlation mathematically or computationally is to sample the two waveforms as shown here, multiply the sample values (a_n x b_n etc.) and add the products together.

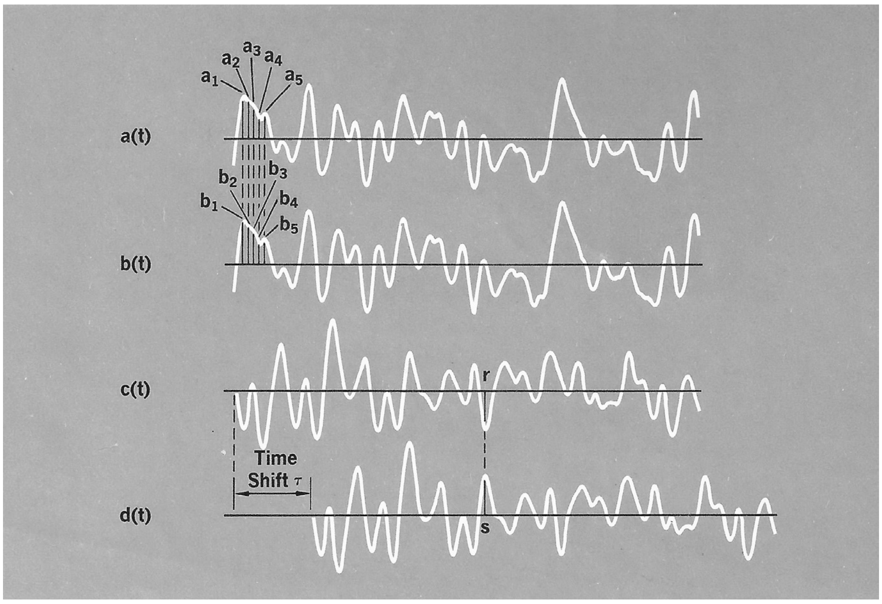

If the waveforms are very similar, the positive and negative samples will both produce positive multiplication products, so the sum will be large. If they differ, then both positive and negative products occur and the sum is much smaller. If the signals are completely uncorrelated, eventually the sum will average to zero.

In the lower two waveforms from the illustration above, the signals are identical but time shifted, by an amount say T. Sampling both waveforms at the same point in time will yield low correlation, but if the lower one is sampled T seconds later then there will be high correlation.

There are two types of correlation measurement: Autocorrelation and Crosscorrelation. With autocorrelation we are simply measuring a single variable and comparing the correlation between the current value and a sample taken some time later, say T seconds. As T is varied, the result is the autocorrelation function of the signal. Periodic signals, such as a sine wave, will generate a repetitive autocorrelation function with correlation peaks at the repeat points in the signal. At the other extreme, truly random noise never repeats, so the autocorrelation function has only one narrow peak at T=0. The width of the peak is dependant on the bandwidth of the noise, wide bandwidth creating a very narrow autocorrelation function in the time domain. In fact this property of wideband noise (or white noise) creates the equivalent of an impulse in the time domain which as we will see later is interesting for characterising the response of systems.

While autocorrelation measures the property of a signal, crosscorrelation measures the property of a system since it evaluates the correlation between two signals, for example between the input test signal and the output system response. From this can be deduced the overall system response.

Impulse Response

When you strike a bell with a hammer it rings. The hammer delivers an impulse and excites the natural response of the bell, which is its impulse response. This analogy applies to any mechanical, electrical or control system. The impulse response gives information on the delay before the system responds and whether it under-shoots or over-shoots with the correction from its control loop. In fact the impulse response contains all the information needed to characterise a linear system and from it one can describe how a system will respond to any arbitrary input or disturbance.

As with the bell, we could measure the impulse response of the system by exciting it with an impulse or train of impulses and observe the output on an oscilloscope. However, impulses are dangerous and may cause overload or saturation which would be undesirable in a working system. Of course, small impulses could be used, but if they were small enough to be safe, they usually produce outputs which are so small that they are obscured by background disturbances.

As already mentioned, the autocorrelation function for random noise is an impulse and it can be shown mathematically that the crosscorrelation function between the output and input of a system is proportional to the system's impulse response. So by using random noise as a stimulus, it is possible to derive the impulse response without using a big hammer! To get useful results, the noise stimulus needs to be of much wider bandwidth than that of the system under test (to create a suitable autocorrelation impulse) and the crosscorrelation process will require averaging to remove uncorrelated system noise.

The mathematical background to this method is beyond the scope of this present chapter, but two useful articles on the theory written by the South Queensferry team were published in the *HP journal* – see the references below[2].

Random and Pseudo-Random Noise

Having established the advantages of using random noise as a test stimulus, the next problem is how to generate it. Truly random noise is usually generated by a "natural" source such as gas-discharge tubes, temperature-limited diodes or by amplifying thermal noise from a resistor. The problem with natural sources is that the long-term statistics are unknown or variable. This is particularly true at low audio or sub-audio frequencies, which is exactly the range of interest in systems and mechanical/structural analysis.

[2] "Testing with Pseudo-Random and Random Noise", HP Journal September 1967, pp. 18-20
http://www.hparchive.com/Journals/HPJ-1967-09.pdf
"Correlation, Signal Averaging, and Probability Analysis" HP Journal November 1969 pp. 2-8
http://www.hparchive.com/Journals/HPJ-1969-11.pdf

The power available from natural sources at these very low bandwidths is small and so must be amplified by a direct-coupled (D.C.) amplifier which may well add drift and internal noise to the test stimulus. Although truly random noise might seem desirable to produce a singular autocorrelation function (an impulse), it means that no two sets of measurements taken over a finite measurement time will ever be exactly the same. This inherent variance is problem when making comparative measurements.

All these limitations led to an interest in the late 1950s and early 60s in synthesizing noise using digital techniques rather than natural sources, and the advent of transistor electronics made the idea a practical proposition. As Gordon Roberts pointed out, "*Digital techniques offer the ability to generate waveforms of known amplitude distribution, and with known spectral properties. All we need is a stream of random numbers.*" While the digitally synthesized noise appears random to the system under test, it is in fact deterministic (pseudo-random) and has completely predictable statistics so that successive measurements can be made under reliably identical conditions.

"*One of the attractive features of low-level pseudo-random noise is that you can hide it within the natural background noise of a physical system",* Gordon Roberts commented. *"Regardless of how small your pseudo-random noise is, you, and only you, will be able to detect it at various points within the system. This feature is valuable if larger disturbances might cause hazardous responses (e.g. in a nuclear reactor or weapon system) or make a hidden message too visible to an unauthorised observer. This is the basis of some work done by Brian Finnie for his PhD at Edinburgh University. The intended application was to measure the frequency response of a nuclear reactor. We now know that the properties of pseudo-random signals had been used before, but were still Official Secrets at the time when we were developing these techniques and products. One application was in code breaking at Bletchley Park.*"

The core of digital noise synthesis is a Pseudo Random Binary Sequence (PRBS) generator:

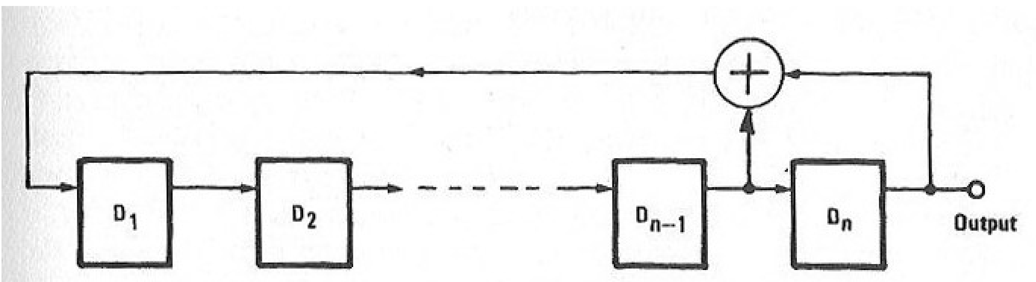

The PRBS generator is a feedback shift-register where two or more outputs from various stages are fed back to the input via Exclusive-OR gates. This defines the run properties of the PRBS, and of particular interest is the family of "maximal-length" sequences. The length of the maximal-length sequence before it repeats is given by the equation 2^N-1, where N is the number of shift register stages.

A shift-register stage is a special-purpose flip-flop or bi-stable which acts as a store for one binary bit of data ('0' or '1'). The individual shift-register stages are connected in cascade so that, on receipt of a shift or clock pulse, the information they contain is stepped progressively along the chain, as if on a conveyor belt. Since it is a feedback shift-register, the first stage receives its input from the Exclusive-OR gates.

Once started, the PRBS generator will produce the same pseudo-random sequence repeatedly at a rate of the clock frequency divided by the sequence length ($f_c/(2^N-1)$), Alternatively, the period of the PRBS is T seconds where T equals the clock period (t_c) multiplied by the sequence length ($T = t_c \times (2^N-1)$).

The output of the PRBS generator is a sequence of binary ones and zeroes. This is useful in some applications, for example when testing a system with an on/off control function such as a solenoid-operated valve. Another important application is the testing of a digital transmission system when the PRBS is used to simulate random live traffic. In Chapter 9, we looked at a whole family of Queensferry instruments, developed from the 1970s onwards, used for checking error performance based on this principle.

For most applications envisaged by Gordon Roberts and Brian Finnie, the need was for an analogue Gaussian noise signal similar to that produced by a natural source. To obtain this, the PRBS is fed through a low-pass filter which changes the binary sequence into a multi-level signal through the filter's own energy storage. To work effectively, the low-pass filter should have a cut-off frequency approximately 1/20th of the clock frequency used in the shift register.

The statistics of this multi-level signal are defined by the Probability Density Function (PDF), which for truly random Gaussian noise is this bell-shaped curve:

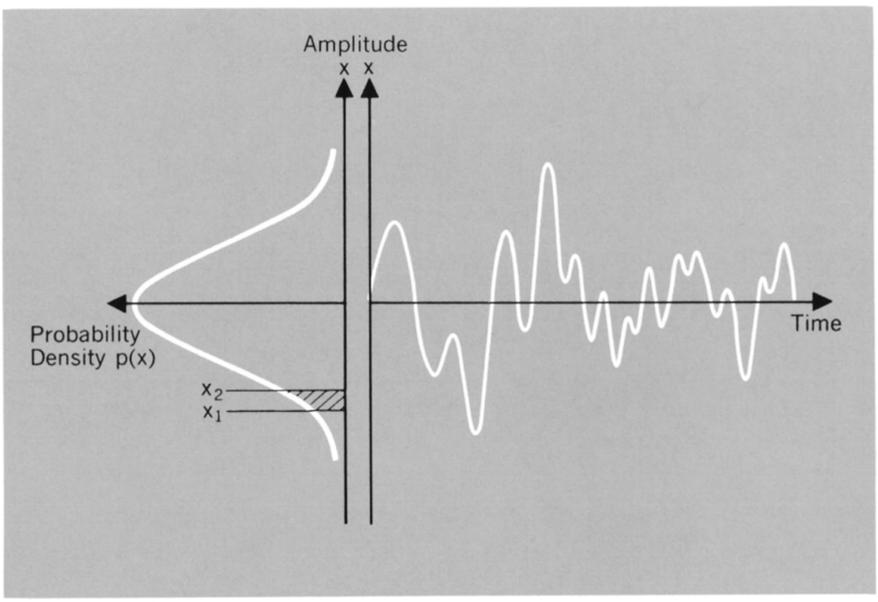

The probability curve shows the likelihood that a certain amplitude will be reached: large amplitude peaks having a low probability. The positive and negative tails of the distribution are infinite for ideal random noise, meaning that infinite amplitude is theoretically possible if you wait long enough! Synthesized noise generated by a PRBS has a similar PDF, however the maximum possible value is limited and determined by the sequence. The longer the PRBS sequence (larger N), the closer the synthesized noise PDF gets to the perfect Gaussian distribution. This is because the longer sequences create longer runs of "ones" and "zeroes" which give rise to higher analogue noise voltage peaks after filtering. Sequences of greater than 8191 (N>13) have very good statistical properties.

This oscilloscope picture illustrates the properties of pseudorandom Gaussian noise. While at first glance the trace looks random, it does in fact repeat itself a couple of times across the screen, and the oscilloscope is able to obtain a stable trigger which it would never do with truly random noise. This is the analogue noise output of the 3722A Noise Generator described later. In this case the sequence length is 1023 or N=10.

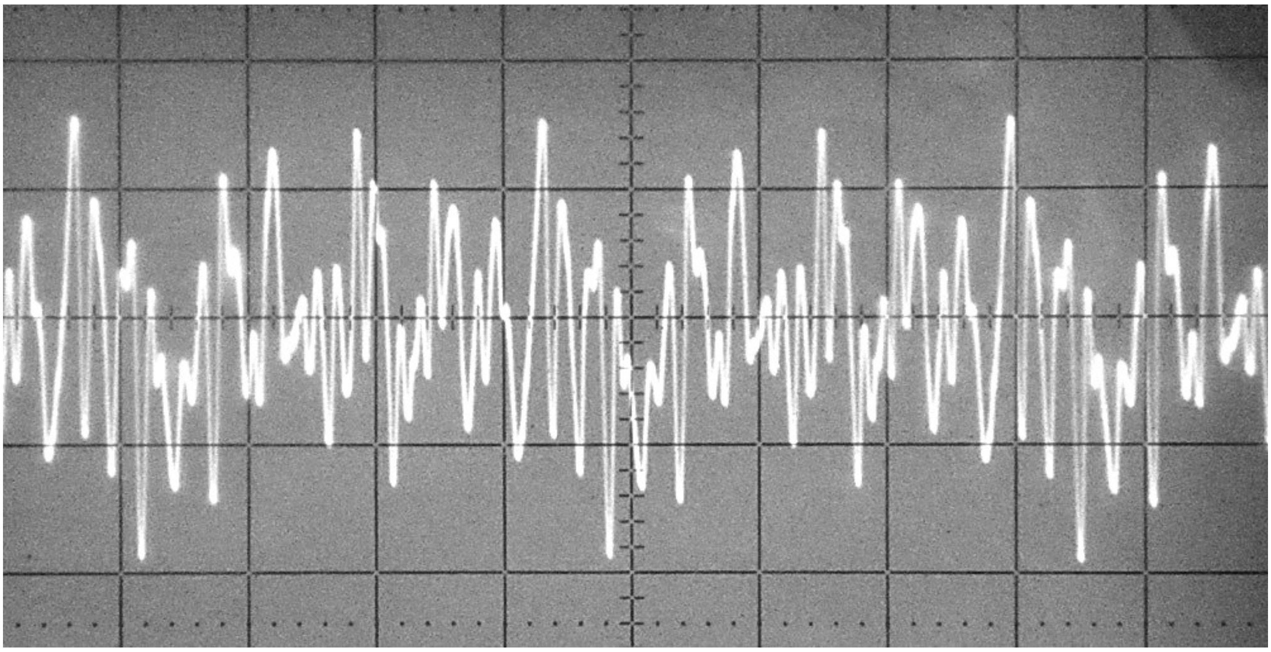

The autocorrelation function of pseudo random noise has the desired "impulse" characteristics, similar to true random noise. However, while true random noise has a single autocorrelation peak at t=0, the synthesized noise has an autocorrelation peak every time the sequence repeats, that is every T seconds as one can see from the photo above. This is very significant since if the measurement time is also made equal to T seconds, the results of an experiment will be identical on every repetition as long as nothing else has changed. This eliminates the variance incurred with true random noise and also allows successive measurements to average out the uncorrelated background noise.

This simplified description of correlation and digital noise synthesis will hopefully explain the background behind ideas Gordon Roberts and Brian Finnie presented to HP in 1965. They proposed two instruments that were completely new inside HP, and indeed anywhere else, at the time. These were firstly a digital noise generator suitable as a stimulus for system testing, and secondly a real-time correlator that would for the first time directly display autocorrelation and crosscorrelation measurements made using digital noise stimulus.

In the following pages, we will look in more detail at the three products that came out of this development during the period from 1965 to 1972. At South Queensferry, this family was referred to as Dynamic Signal Analysis, or DSA for short.

3722A Noise Generator (Unit Preserved)

Development work soon started on the new 3722A Noise Generator in 1965. At the time construction of the new HP facility at South Queensferry had just begun and would not be completed until summer 1966, so the design team had to work elsewhere. While Peter Carmichael and Finlay Mackenzie working on the Microwave Link Analyzer used the Burgh Chambers in South Queensferry village, Gordon Roberts and his team worked in a room at his Edinburgh home.

The 3722A is basically quite a simple instrument consisting mainly of the digital PRBS generator, described earlier, with filtering at the output to create quasi-Gaussian noise. The PRBS shift register was driven by a clock signal derived from a 3 MHz crystal oscillator so that the frequency parameters were very accurately controlled. The clock rate could be varied over a wide frequency range so that the noise signal could be optimised for many different applications from very slow-acting control systems all the way up to high audio frequencies and beyond. This was done by dividing the crystal oscillator output through a large number of decade digital dividers: the shortest clock period was one microsecond (1 MHz clock rate) all the way to 333 seconds or 5.5 minutes (0.00015 Hz)!

The design team needed a lot of digital building blocks including dividers and shift register stages. In the mid 1960s, the first families of logic integrated circuits (ICs) had started to be introduced by companies like Fairchild, Motorola and Texas Instruments. These might have been a possibility for the new instrument, however the designers decided to stick with discrete transistor logic circuits as these provided more flexibility at the time. Similar building blocks had already been designed for HP's range of digital frequency counters in the early 1960s, and in fact the first two products developed by HP in the UK were 2 and 10 MHz frequency counters, the 3734A and 3735A. As the objective was to get the 3722A into the market as quickly as possible, they used some tried and tested designs.

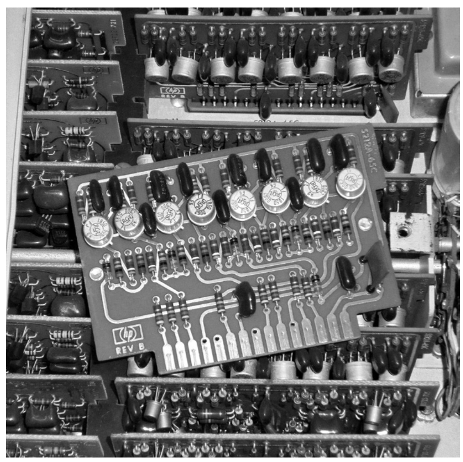

The decade dividers needed for the clock source were already available as single board modules (one is shown here) from the 5212A 300 kHz frequency counter developed by Frequency and Time Division at Santa Clara in the early 1960s. Seven of these modules were used in the new design, and in the preserved unit they are clearly imports as the PC board material is different and they still have the 5212 part number. These decade divider boards had four two-transistor bi-stables or flip-flops with feedback. The boards were roughly 4" x 2", quite small, and provided a sort of template for the rest of the design. Many of the Queensferry-designed boards used a similar arrangement, with four bi-stables per board, for example the higher-speed dividers and also the important shift-registers, which also had four stages per board. Over 40 of these small PCB assemblies were used in the final product – a couple of years later, many of these boards could have been replaced by single logic ICs.

The central part of the design was of course the feedback shift register which generated the PRBS. The 3722A had up to 20 stages in the shift register for PRBS generation. There were an additional 12 stages beyond that in the shift register, making 32 in all, and we'll look at the purpose of these shortly. As mentioned earlier, the length of the PRBS was determined by the number of shift register stages (N) connected by the Exclusive-OR feedback loop to generate a maximal length sequence (2^N-1). The 3722A allowed the user to set N between 4 (sequence length 15 bits) up to 20 which gave a sequence length of over one million which for most practical applications was virtually the same as random noise. If you wanted truly random noise, you set the sequence switch to INFINITE, which disconnected the feedback loop and fed the shift register input from an analogue noise source – amplified transistor noise.

The clock period could be set from 1 microsecond to 333 seconds in 18 steps. With the shortest sequence and clock period, the sequence repeated every 15 microseconds. At the other extreme, the longest sequence (N=20) with the 333 second clock period gave a repeat every 11 years!

As discussed in the introduction, to obtain a Gaussian analogue noise signal, the PRBS from the shift register needed to be passed through a low-pass filter with a cut-off frequency 1/20th of the clock frequency. This was not a problem at the higher clock frequencies (at one microsecond period, the cut-off was 50 kHz), but at lower rates when the required cut-off was below 1 Hz, the analogue solution was problematic in terms of stability and settling time.

The designers found a brilliant solution to this problem, perhaps tapping into their academic knowledge of the mathematics of impulse response and Fourier transforms. Their answer to the low-frequency filter problem was a transversal digital filter. It was the cleverest part of the new instrument.

The theory of the transversal filter is that it synthesizes the time-domain impulse response of the desired filter, rather than its frequency response which is the conventional approach using analogue circuitry. Since the circuit is working in the time domain, multiple successively delayed replicas of the input signal are summed together through appropriately dimensioned weighting networks, meaning that different amounts of the various delayed samples are added together to produce the composite output. The weighting coefficients represent the impulse response of the desired filter: the greater the number of delayed samples, the more accurate the approximation to the ideal filter response.

The ideal low-pass filter has an equivalent sinx/x impulse response in the time domain (the Fourier transform):

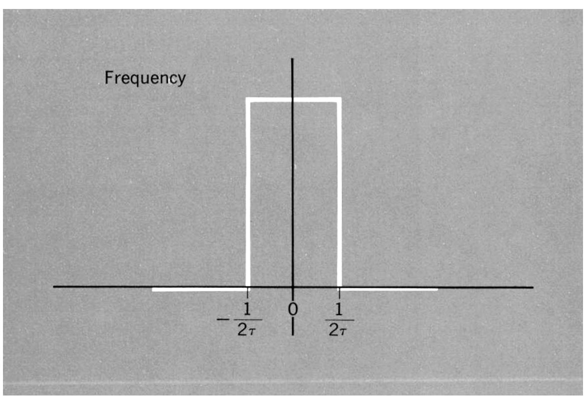

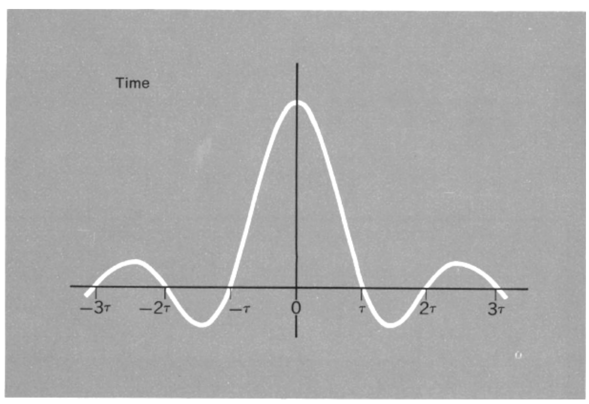

The beauty of the 3722A design was that the first 20 stages of the "delay line" were already there as part of the PRBS generator, with the PRBS values delayed by one clock period at each successive stage of the shift-register. To get a satisfactory impulse response and corresponding low-pass characteristics, the designers added a further 12 stages to the shift-register, giving 32 delay elements all together. The "weighting" required at each stage for the transversal filter was simply achieved using resistors graded to follow an approximation to the sinx/x curve, all being combined in a current-summing amplifier.

Brian Finnie recalled how the filter design weighting coefficients were optimised:

"*We used the then state-of-the-art Atlas Computer*[3] *at Manchester University for the design of the digital filter to achieve the sharpest frequency roll off with the minimum number of (expensive) shift register stages.*"

As shown in the above diagrams, the cut-off frequency for the low-pass filter is related to the first null in the impulse response. The filter was designed so that the -3db cut-off frequency was 1/20th of the clock frequency. This meant that the low-pass filter scaled automatically to the rate at which the PRBS was being clocked. It was a very elegant solution that would naturally work down to arbitrarily low frequencies and with complete stability.

The only consideration was that because the filter had only 32 stages, there was some quantization of the output waveform, giving it a stepped appearance. Given the speed of the logic circuits, these steps were very fast and the designers considered that this might interfere with some higher frequency test setups. So, analogue smoothing or filtering was added to suppress these steps on the higher frequency ranges. Very slow systems would be unaffected by the high-frequency components.

In the 21st century digital age, digital filters and signal processing are everywhere, but in 1965, this must have been one of the earliest applications of the idea and almost certainly the first in HP.

The 3722A provided both analogue (Gaussian) and digital outputs, and had a number of additional features such as a synchronizing pulse for the start of a PRBS sequence, the facility for running a specific number of sequence repeats, as well as being able to "RESET", "RUN" and "HOLD" a PRBS. These features could be programmed via a 36-way "Amphenol" connector on the back panel, making the instrument one of the earliest with computer programme control. An interesting observation is that with its digital PRBS output, it was South Queensferry's first digital pattern generator, a field that became the Division's mainstream business in later years, although it would be another six year before the first true telecom pattern generator and error detector emerged, the 3760A/61A – see Chapter 9.

The product now looks antiquated with its front panel featuring three large rotary switches (sequence length, clock, and output level) and the transistor logic inside. But this conceals its innovation, which for the time was impressive. It was probably the first time that

[3] The Ferranti Atlas was first installed at Manchester and commissioned in December 1962. It was one of the world's first super-computers. No doubt Gordon and Brian had access to this computer through their previous association with Manchester University.

complex analogue test signals had been generated and filtered entirely digitally. This gave the new instrument some important advantages:

- Predictable statistics and repeatability
- Accurately defined frequency characteristics
- Stable and calibrated output levels even at very low frequencies

Its significance was recognised, as the Noise Generator had an entire issue of the *HP Journal* devoted to it in September 1967[4]. Written by the South Queensferry design team, it has a lot more information on the development of the 3722A and the mathematics behind it.

Brian Finnie did the initial design work on the Noise Generator and in 1966 George Anderson took over as Project Leader when he joined HP from the Edinburgh Royal Observatory. Meanwhile, Brian went on to develop the second new instrument, the 3721A Correlator described next. Gordon Roberts became Technical Manager at the new factory, and R&D Manager in 1968.

George Anderson, Brian Finnie and Gordon Roberts (left to right), with the Noise Generator

The 3722A Noise Generator was introduced sometime in mid-1967, slightly ahead of the Microwave Link Analyzer, so it was probably the first product to be designed and built entirely at the new facility. David Simpson, General Manager, commented in the Division newsletter, "*The 3722 was shown at the large electronics show in San Francisco, Wescon, and made quite a technical impact, which is reflected in a forecast of 15 units/month for 1968.*" It certainly lived up to this confident prediction, with sales of 257 during that year. By the early 1970s, it had settled down to around 12 units/month. At launch it cost $2.65k rising to $4.2k in the late 1970s. The Noise Generator was in production until 1980 and a total of 1590 units were sold generating about $5M of business.

Finally, the preserved unit deserves some comment as it is quite a survivor. It has serial number U717 00057, indicating it was built in

4 http://www.hparchive.com/Journals/HPJ-1967-09.pdf

April/May 1967 and was a production prototype. It was almost certainly used by some of the engineers mentioned in this chapter. Over the years it obviously found use as a general purpose noise generator at the factory, surviving numerous moves and changes of strategy. It was discovered in a stack of surplus test equipment as the factory finally closed down in the summer of 2010, 43 years after it was built. It had spanned virtually the whole life of the site, being one of the first units designed and built there, and then witnessed the decades of growth and the catastrophic decline that followed. It still worked, a testament to its robust digital implementation.

The 3721A Correlator (Unit Preserved)

A real-time correlator was the ultimate goal for all this development work. In the opening section of this chapter we saw how the use of autocorrelation and crosscorrelation provided a valuable way of characterising systems and signals. Gordon Roberts and Brian Finnie had done research at Edinburgh University on this topic, and Brian had done his PhD specifically on real-time correlation, so he started an investigation into the new product in 1966. The 3722A Noise Generator provided the digitally synthesized test signal, while the correlator would provide the matching measurement instrument.

Correlation techniques had been in use for some time as a means of extracting coherent signals from noise and for analysis of system performance. As we saw earlier in this chapter, the method requires sampling a large amount of data and then computing autocorrelation or crosscorrelation functions by multiplication and addition.

One method of doing this was to record the data and then process it later on a computer which could do the number crunching. From the late 1950s, computers were available which could do this quite quickly, but the whole process was laborious and time-consuming – and somewhat impractical if you were using the technique for making adjustments or modifications.

Another possible method was to sample the signals with analogue circuitry and store the sample values in capacitors, making the correlation calculations with analogue computers. This could give faster results, but capacitive storage and averaging became a problem with low frequencies and slow systems because the circuits weren't stable enough over time. Just as the Noise Generator had solved the low frequency problems using digital techniques, the answer was to do the same in the real-time correlator. So far, so good.

The next major challenge was how to implement digital circuits that would provide storage, multiplication and averaging, and still fit into a portable measuring instrument. By the middle to late 1960s, the first digital integrated circuits had begun to appear, but they had relatively low functionality (simple logic gates or bi-stable flip-flops).

HP's first minicomputer (the 2116A) appeared in 1966, but these were pretty large pieces of kit despite using the first generation ICs. The correlator needed to be more compact than this, and needed to include a CRT display for presenting the real-time results.

Like the designers and software writers of the first generation microprocessor systems in the early 1970s, a great deal of ingenuity and invention was needed to work around the severe limitations of that early digital hardware. In fact, if early microprocessor had been around, it would have made the correlator development a lot easier, but they weren't in 1966. So Brian Finnie and the team had to devise a way of making a special purpose correlation computer with the logic ICs available.

To understand in broad terms how they did this, it is best to examine the block diagram of the correlator they devised, shown here.

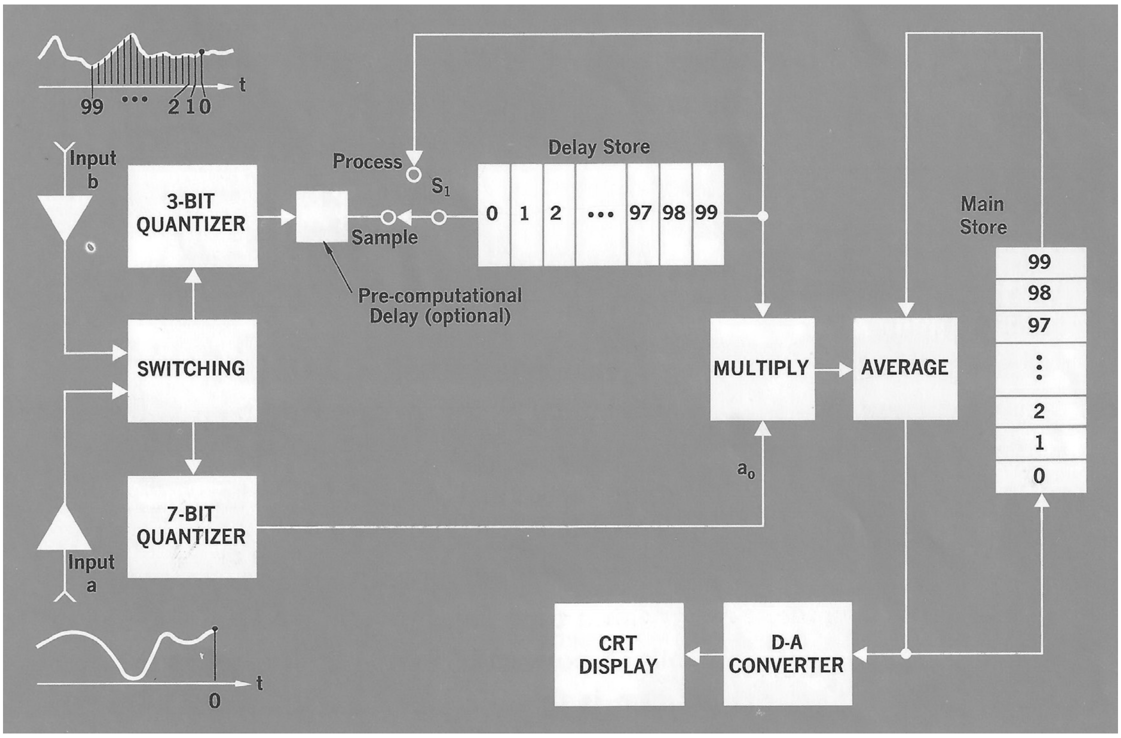

The instrument had two inputs so that crosscorrelation measurements could be made. For autocorrelation measurements a single input was used, and the correlation function computed between the input signal and a delayed version of it. The two inputs had analogue range-switching amplifiers similar to an oscilloscope which adjusted the input signal level to be optimum for the analogue-to-digital converters (ADCs), or quantizers.

The correlator took 100 samples of the input signal, which were displayed as 100 points across the CRT screen. Thus the horizontal time-scale on the screen would be 100 times the sampling interval, which could be set between one microsecond and one second, depending on the received signal bandwidth.

To make measurements, the correlator needed to store/delay up to 100 samples of data, multiply two binary words and average the result. Using the 7-bit quantizer, each of the 100 samples would be a 7-bit word, which would need to be stored in a shift register. With the devices available in the late 1960s, this would require a colossal number of digital ICs: 350 dual D-type Flip-Flops plus ancillary logic. On top of that, multiplying two 7-bit binary numbers would require a substantial amount of logic circuitry, all consuming large amounts of power. Two ingenious ideas overcame these difficult problems of scale in a portable instrument.

The first inventive step was to explore how many bits of quantization were actually required to achieve acceptable measurement accuracy. While the reference channel could use a 7-bit ADC as it was fed directly to the multiplier, the problem was the delayed channel as it required 100 shift-register stages. The designers hit on the idea of a much coarser 3-bit encoder for the delayed channel which would operate faster and require fewer ICs for the 100 stage shift-register. The downside was that it would produce large quantization steps, referred to as “quantization noise” since it is an error between the analogue signal and the digital representation. Computer analysis, again using Manchester University’s Atlas Computer, showed that the 3-bit converter would provide satisfactory results because the correlation measurement involves averaging which would reduce the quantization noise. The performance could be improved further by adding Gaussian noise to the input signal (referred to as “dither”) to randomise the quantization.

Another advantage of the simple 3-bit encoder was that the output values were defined as simple binary values (±1, ±2, or ±4). This meant that multiplication with another binary word (the output of the 7-bit converter), only involved shift operations saving hardware in the multiplier.

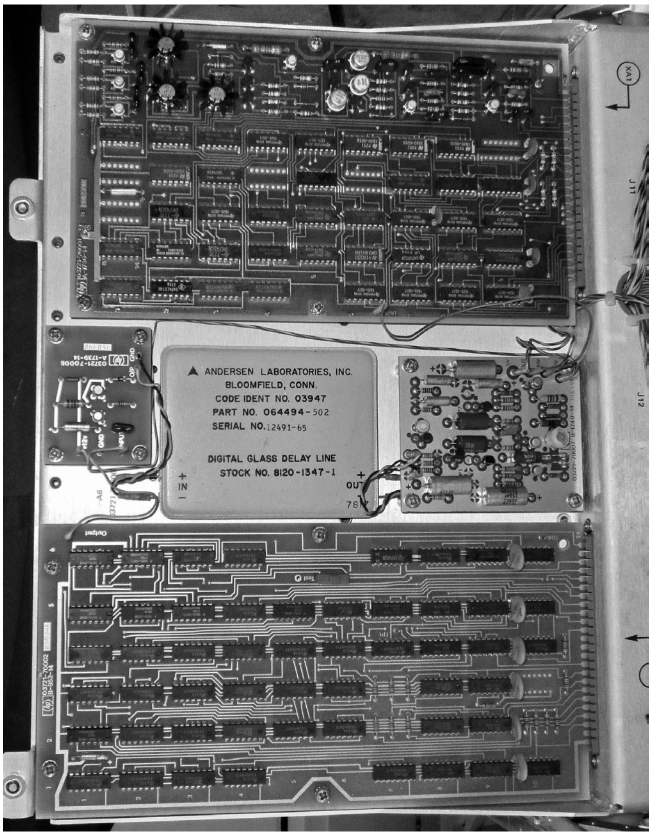

The second innovation was the implementation of the main store. Here it was necessary to store 100 computed results from the multiplier and accumulate multiple samples to provide the averaging. This needed to be 24 bits wide (2400 bits of stored data), well beyond a practical implementation using logic ICs. The solution here was to use an ultrasonic delay line to provide a pipeline store. Clocked serially at a rate of 18 MHz, it had a complete cycle[5] of 136 microseconds. It was a compact way of storing the data, and the access delay caused by a serial store was not a problem in the correlator since the information was being processed and displayed serially anyway. This photo shows two of the five main processor boards on a hinged panel, with the glass delay line in the centre.

Output from the main memory to the display was via a shift register which “windowed” eight bits to a buffer and then to a digital-to-analogue converter to provide the vertical deflection on the CRT. Selecting which eight of the 24 possible bits in the stored word depended on the gain settings for the display. In fact the absolute calibration of the correlator was a complex combination of analogue input scaling and the digital settings, so this was calculated and displayed automatically by annunciators to the left of the CRT on the front panel, simplifying operation.

[5] The actual store had 102 words, two additional words being used for housekeeping and control. The glass delay line was a lump of glass with multiple reflecting surfaces so the wave would bounce round inside to provide the required delay. Early units had a transmission loss of about 60dB, though later ones, supplied by a US company which were about 4 inches in diameter, had a loss of 30 to 40dB.

Apart from the two correlation measurements, the digital hardware could be configured for doing signal recovery (extracting a coherent signal buried in random noise) which was in effect a correlation type measurement. It could also analyse the probability density for a received signal (such as the Gaussian curve shown earlier). The probability density function was a useful measurement for detecting non-linearity in the system under test since non-linearity would alter the theoretical distribution (for example compressing the peaks of the signal or tails of the distribution).

Brian Finnie recalled that the team used the newly introduced HP 2116A computer to optimise the Correlator design. "*Effectively we used an early version of Computer Aided Design to simulate the system for answering design questions – another first for the South Queensferry division.*"

Much more information about the design of the 3721A Correlator, and its various measurements, was published in the *HP Journal* of November 1969[6], which also describes some fascinating practical examples of correlation in action. These included:

- In-service characterisation of large process plant and mechanical devices such as satellite dishes, working in conjunction with the 3722 Noise Generator.
- Using the Noise Generator as a stimulus for measurement of acoustic absorption coefficients and sound direction.
- Using crosscorrelation to detect vibration in rotating machinery and noise transmission paths, and contactless velocity measurement in a strip-rolling mill.
- Detecting repetitive signals buried deep in noise such as very weak radio sources in space received by high-gain radio telescopes.

Here is a typical display from the 3721A Correlator, published in the product data sheet from August 1969. The applications for the new Correlator were wide and varied:

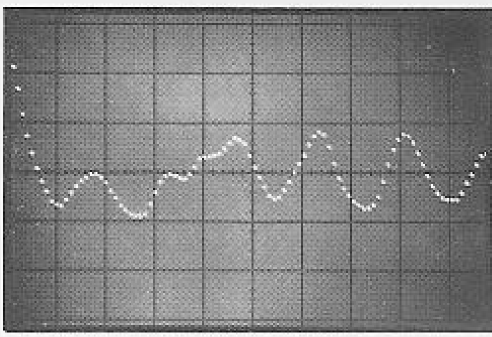

Autocorrelation of background noise in office area. Shows periodic component due to ventilation fan, normally inaudible. Timescale 1mS/mm.

For the late 1960s, the 3721A Correlator was a remarkably innovative product and probably the first HP instruments to make measurements entirely using digital signal processing. With the technology available, this was a challenging project and there were a lot of large boards full of logic ICs. The service manual for the Correlator had to be published in two volumes!

[6] "A Calibrated Real-Time Correlator/Averager/Probability Analyzer" by George Anderson and Michael Perry, HP Journal, November 1969, pp 9-20.
http://www.hparchive.com/Journals/HPJ-1969-11.pdf

The full digital implementation made it an excellent candidate for early computer-controlled measurement systems using HP's first computer, the 2116. The 100 displayed measurement points could be fed to a computer as parallel 14-bit binary words via the rear-panel "Amphenol" connector. This also provided remote control of the RUN, HOLD and RESET commands in the instrument, similar to the functions on the 3722A Noise Generator. Thus a complete automatic measurement system was possible. Computer software could convert the time-domain data from the Correlator (using the Fourier Transform) into spectrum analysis and system transfer function.

As with the Noise Generator, George Anderson took over as project leader after the initial development work by Brian Finnie. A team of around eight electronic and mechanical designers worked on the product for two to three years, and it was introduced during the summer of 1969. It was quite a large instrument, 11" high, 17" wide and 19" deep with a rectangular centre-mounted CRT screen. At launch, the list price was $8.5k, rising to $11.2k in 1979. It was in production for about ten years and a total of 610 units were sold, bringing in about $5.8M.

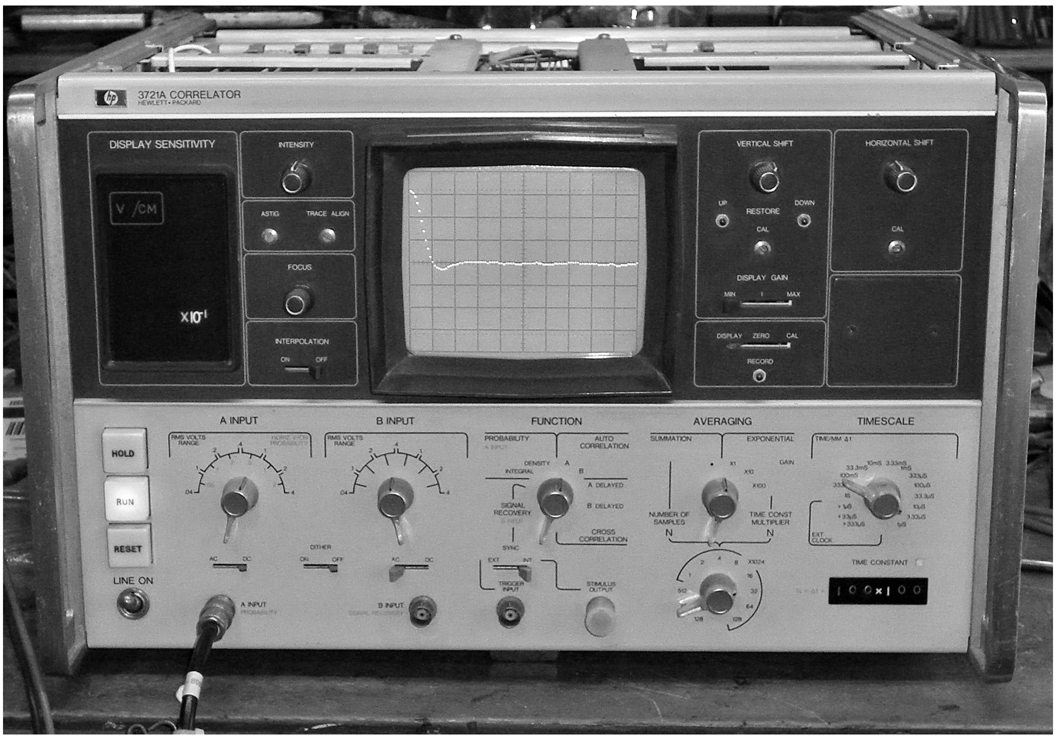

The preserved 3721A Correlator displaying the autocorrelation function of the 3722A Noise Generator

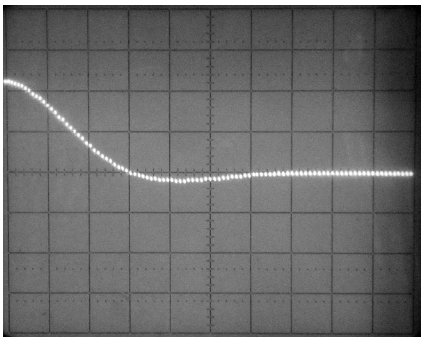

Autocorrelation of noise

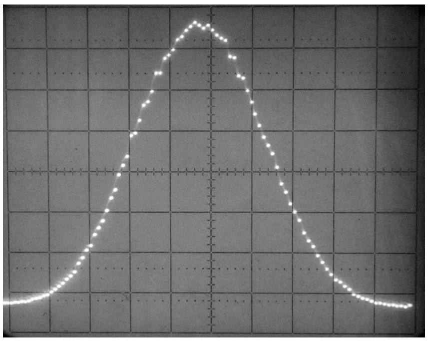

PDF of noise

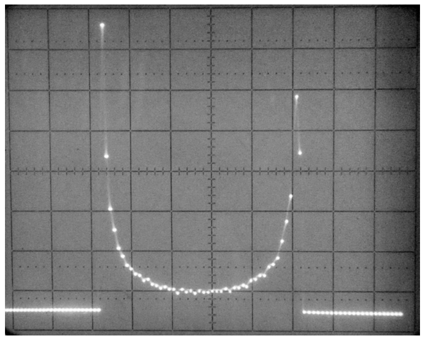

PDF of sinewave

3720A Spectrum Display

As mentioned earlier, there is a duality between the time-domain impulse response and the frequency domain spectral response, through the mathematical equations of the Fourier Transform. The 3721A Correlator produced time domain measurement data, or more accurately discrete time domain information at 100 sampling points across the measurement period. This data could be transferred to a digital computer and converted to frequency domain information by using a Discrete Fourier Transform (DFT)[7]. The mathematical equations for this conversion into sets of real and imaginary coefficients, assuming 100 sample points, are:

$$R(n\Delta f) = \frac{1}{100} \sum_{k=0}^{99} x(k\Delta t).\cos(2\pi n\Delta f.k\Delta t)$$

$$I(n\Delta f) = \frac{1}{100} \sum_{k=0}^{99} x(k\Delta t).\sin(2\pi n\Delta f.k\Delta t)$$

where n = 1, 2, 3, etc., and $\Delta f = 1/100\Delta t$ where Δt is the sampling time interval used for the measurement in the Correlator. Each real or imaginary coefficient at frequency $n\Delta f$ is computed by summing the 100 (k = 0 to 99) products from the above equations for the 100 sets of time-domain data ($x(k\Delta t)$) received from the Correlator. This numerical calculation can be done quite readily on a digital computer, and the South Queensferry team developed software on the 2116B computer to do just that.

While the measurements from the 3721A Correlator were all in the time domain, much engineering and scientific work uses frequency domain measurements such as power spectra and transfer functions. Intuitively, it is often easier it to associate frequency domain measurements with physical phenomena rather than the impulse response.
The computer software solution provided this information but it was a cumbersome set-up for practical measurements, particularly when the tests were used for making real-time adjustments. Around 1969, the Correlator design team had the idea for a companion instrument that would take the time-domain measurement results and convert them to the frequency domain in a user-friendly way. This became the 3720A Spectrum Display.

It wasn't a standalone instrument, rather it was an accessory to the 3721A Correlator, in effect a dedicated Fourier Transform computer. As with the two earlier instruments, the challenge for the team was how to build it with the very limited digital hardware of the time and fit it into a portable instrument with a display. No point in building another 2116 computer! David Dack, (later to become famous in the annals of HP South Queensferry) was deeply involved in theoretical studies and preliminary definition of the new instrument.

As shown above, the Discrete Fourier Transform (DFT) involves complex multiplication and the summing of a large number of results. In 1969, this was a bit too much for a full digital implementation, so the 3720A uses a hybrid approach of digital and analogue processing – mostly analogue, as shown in the block diagram on the next page.

[7] It is discrete since it has a finite number of data points, suitable for numeric computing, rather than a continuous mathematical function.

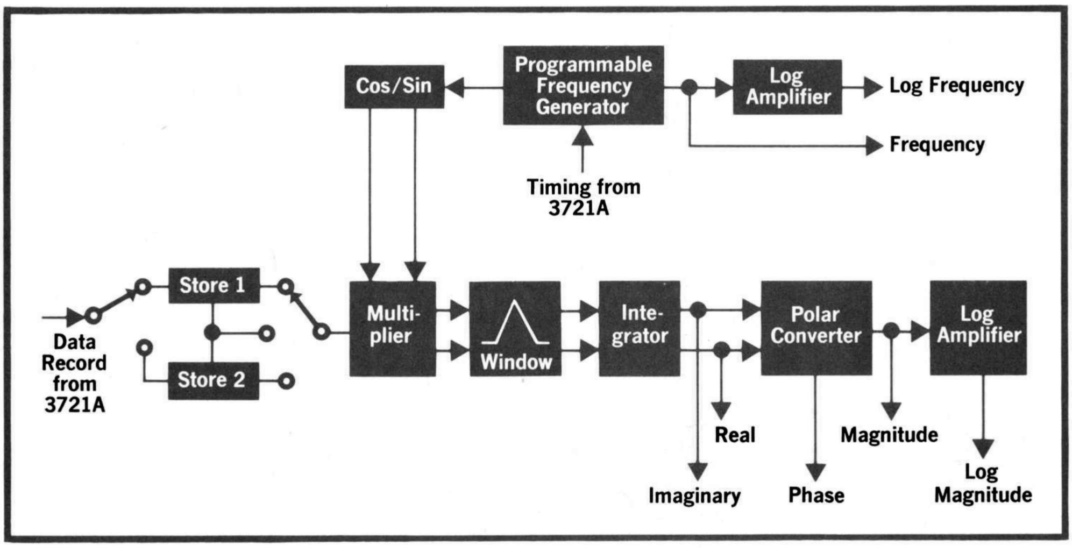

The main digital part of the instrument was the two data stores that took the time-domain records from the Correlator (100 x 12-bit words) via a multi-way cable. To execute the DFT equations, cosine and sine waves of controlled frequency were required. These were produced by programmable function generator circuit (a triangular wave followed by diode shapers)[8]. Each Fourier coefficient was calculated by summation of the results of hybrid multiplication of the digital record, term by term, by successive analogue samples of the cosine and sine waves. The storage of these terms was in an analogue integrator with storage in two capacitors, one for the imaginary part and the other for the real part. The remaining circuits provided logarithmic and polar conversion of the displayed results using analogue circuitry. To display the accumulating frequency domain results, the 3720A used a special CRT from a storage oscilloscope. The photos below show the displays on the 3720A, derived from the Correlator crosscorrelation measurement (top left).

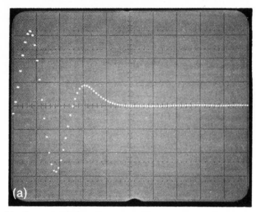

Estimate of the Impulse Response of Band-pass Filter, obtained by Crosscorrelation.

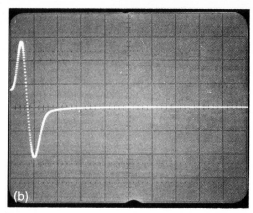

Transform of (a), Real Part versus Frequency.

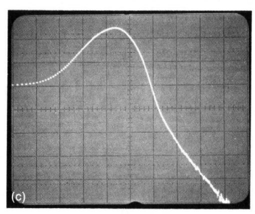

Transform of (a), Log Magnitude versus Log Frequency (Bode diagram).

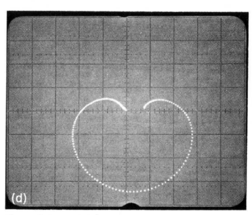

Transform of (a), Real Part versus Imaginary Part (Nyquist diagram).

This is just a brief overview of its operation, however the design team led by David Morrison published a much more detailed description in the November 1972 issue of the *HP Journal*[9]. The 3720A Spectrum Display was introduced during 1972 at a list price of $6k, rising to $8.7k in later years. Around 220 units were shipped between 1972 and 1979, with total sales of about $1.6M.

---

[8] The circuit was based on the recently introduced 3310A Function Generator from Loveland Division

[9] "*Versatile Display Unit Extends Correlator Capability*", by Morrison, Finnie, Patel and Edwards, HP Journal, November 1972 pp. 8-15 http://www.hparchive.com/Journals/HPJ-1972-11.pdf

The DSA Team arrives at a Crossroads

By the early 1970s, South Queensferry division management was faced with a strategic decision as to which direction it would go. Bill Hewlett and Dave Packard encouraged them to focus on one particular market and product area, as other HP divisions tended to do. They believed it was hard to understand the market and compete successfully in more than one area. It looked like the Division would have to choose between the DSA product line and the Microwave Link Analyzer (MLA) together with the other telecom products (such as the new pattern generator and error detector) already under development.

Unfortunately for the DSA team, the Division chose the telecom test direction for two main reasons. Firstly, the 3701A MLA system and its successor the 3710A system proved a spectacular success, with sales far outstripping those of the DSA products. The rapidly growing telecom market and the emergence of digital telephony also looked like it had a lot of future potential. Secondly, from HP's point of view, the Queensferry DSA product development appeared to be duplicating work going on in the Santa Clara Division in California.

In the late 1960s, the Frequency and Time Division at Santa Clara had got interested in the Fast Fourier Transform (FFT), following a theoretical paper (Cooley & Tukey, 1965)[10] which showed how to compute a Discrete Fourier Transform (DFT) much faster than the direct method described earlier in this chapter. The Santa Clara product was basically a digital waveform analyzer, comprising an Analogue to Digital Converter (ADC) and FFT software running on a minicomputer to provide a near real-time spectral display. Their product, the 5450A/51A Fast Fourier Analyzer was introduced around 1970, consisted of a large equipment rack with an HP 2115A computer (later a 2100A), and was aimed at vibration/structural analysis – the same kind of markets that South Queensferry's Correlator addressed. Compared to the Queensferry product combo however, it was very bulky and expensive, as these photos show.

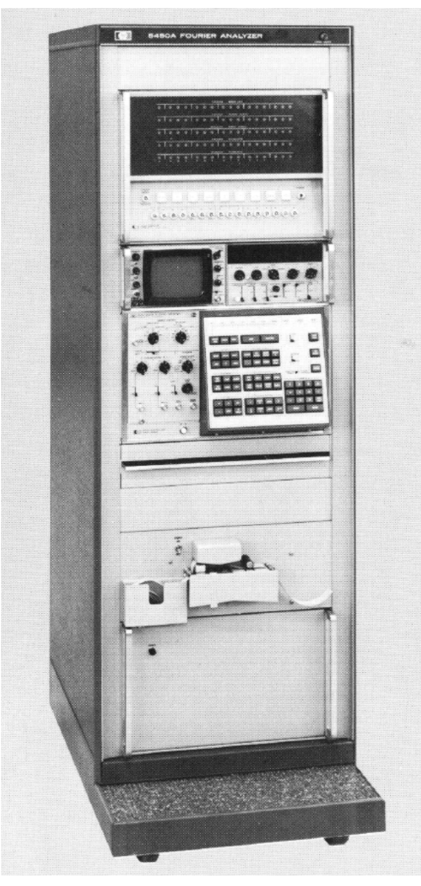

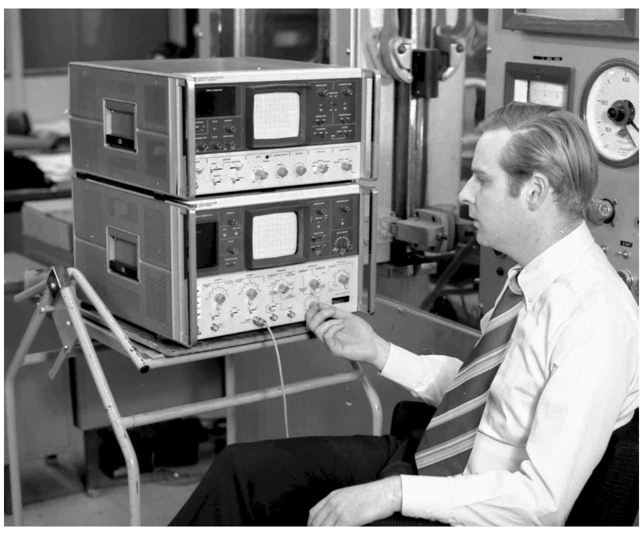

[10] Cooley, James W.; Tukey, John W. (1965). "An algorithm for the machine calculation of complex Fourier series". Mathematics of Computation **19/90**: pp.297–301. http://www.ams.org/journals/mcom/1965-19-090/S0025-5718-1965-0178586-1/S0025-5718-1965-0178586-1.pdf

So, once the 3720A development was finished, the DSA R&D team was dispersed. Some moved on to other projects and to the recently launched Systems Integration Centre. Others decided to look for new opportunities inside and outside HP. In the end relatively few of the key players stayed at South Queensferry. One of the former Correlator team, Dick Watts, moved to Santa Clara and became the marketing manager for FFT products, even creating a specialist sales and application group to sell the systems. It was moderately successful, but it wasn't until the Loveland Division in Colorado introduced its 3582A FFT Spectrum Analyzer (a competitively priced bench instrument) in 1979, that high volume sales developed.

As for the three DSA products, they continued to be marketed throughout the 1970s, but ceased to be of any strategic interest to the Division. They just became a modest but useful source of income. One problem with the products was that their technology quickly became outdated, simply because they were so early into the market. When they were designed in the late 1960s, digital technology was embryonic. By the early 1970s, Intel's first microprocessor entered the market, and there were semiconductor memory chips, while logic ICs had increasingly comprehensive functionality. A second generation version of all three products, introduced by the mid-70s, would have delivered a massive improvement in capability and price/performance. Of course that didn't happen, and the increasingly outdated products soldiered on until 1979. As a result, sales tailed off:

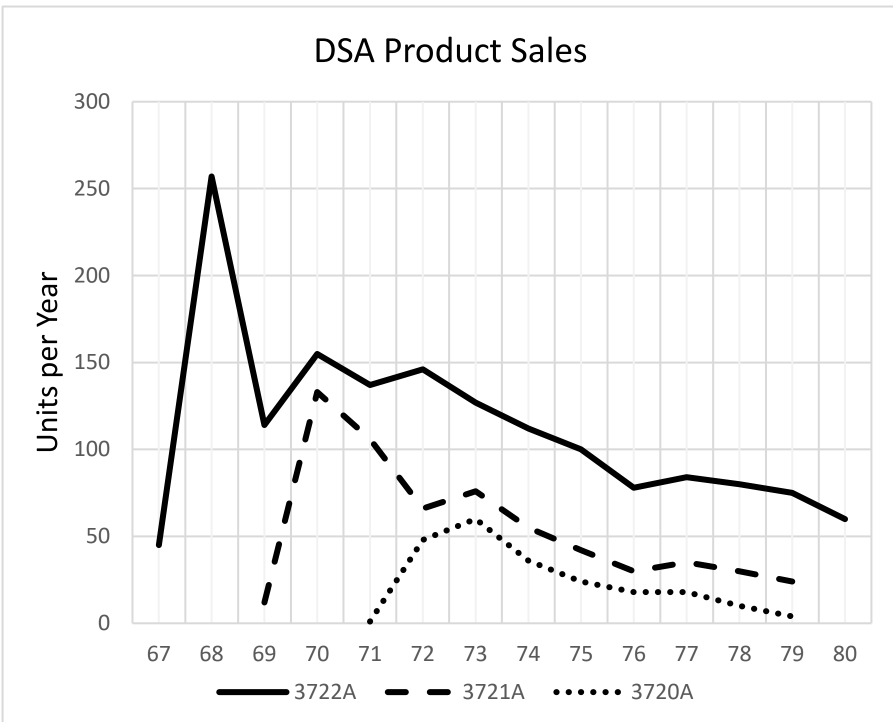

As the years rolled by, the old DSA products were gradually forgotten as everyone concentrated on the telecom business. Younger staff had never even heard of them, but fortunately a Noise Generator and Correlator have survived into preservation.

David Dack, one of the design engineers, commented:

"*The DSA adventure was full of very innovative ideas, but sadly on the wrong track commercially. Many of the measurements could have been done better by a computer-based system, and although computers were expensive, the people who needed this sort of* *analysis could afford them, so there was no substantial market for cheaper alternatives. I am grateful for the experience however, as it convinced me to always go for a digital solution to a problem first and to use analogue design only where essential.*"

In a way it is sad that the brilliant ideas of the "Correlation Pioneers" were never fully developed, but as the Division spread its wings in the 1970s, something had to go. However, the Division did benefit from these products in an oblique way.

The knowledge of PRBS generation accelerated the development of the first pattern generators, and the innovative digital design put the R&D team in a strong position for all the advanced digital signal processing and logic design invented at South Queensferry in the 1970s and 80s. This ended up being the true legacy of the DSA development.

Acknowledgements

Thanks to the following former South Queensferry employees who contributed to this chapter: Gordon Roberts, David Dack and Brian Finnie.

Appendices

Appendix 1

Introduction to HPSQF Product Catalogue

The table in the following pages lists all the instruments and systems designed and built by the Telecom Division at HP South Queensferry (Agilent Technologies from 1999 to 2005). The Telecom Division had various names over the 40 year period, starting with the South Queensferry Division (SQD), then Queensferry Telecom Division (QTD) and in the 1990s, Queensferry Telecom Operation (QTO) and finally Telecom Networks Test Division (TNTD). For many employees, however, it remained QTD despite the name changes.

It was the main division throughout the history of the site, although from the mid-1980s there was also the Queensferry Microwave Operation/Division (QMD) and from around 1994 the Telecom Systems Division (TSD). This catalogue does not include their products. QMD was mainly a manufacturing and re-engineering operation and TSD had one main product, the HP acceSS7 network monitoring system. These divisions and their products are covered in separate chapters and appendices in Volume 1.

The catalogue was mainly compiled using some period documents found at the factory. These include sales records from 1970 onwards, and of particular value, some binders compiled by the Product Support Department over the years. These included the product support plan for each new product and later the obsoletion notices which contained dates and units sold. The catalogue also acts as an index to Volume 2, showing the page where the particular product is mentioned.

More information on the methodology and the meanings of abbreviations can be found at the end of the table.

Please see notes at end of spreadsheet for methodology

QTD/TNTD South Queensferry Product Catalogue (1965 to 2005)

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
CLASSIC PRODUCTS											
3701A	Microwave Link Analyzer System (70 MHz IF)	1967	1973	$8k	1721	$13.8M	Includes 3702A and 3703B	18	NMS T.2010.65.1-3		Peter Carmichael
3701Z	Microwave Link Analyzer System for BELL Std.	1970	1973	$8k	850	$6.8M	Includes 3702Z and 3703Z	21			Finlay Mackenzie
3702A/Z	MLA Receiver	1967	1973				See 3701A/Z Systems		NMS T.2010.65.1-3		Finlay Mackenzie
3702B	MLA Receiver	1972	1985				See 3710A System			HPJ September 1972	Colin Appleyard
3703A/B	Group Delay Detector Plug-in (LF Test Tones)	1967	1985				See 3701A/Z and 3710A Systems		NMS T.2010.65.1-3		
3705A	Group Delay Detector Plug-in (LF & HF Test Tones)	1972	1985				See 3710A System			HPJ September 1972	Norman Edgar
3706A	Group Delay Display (0.1 to 12.4 GHz)						Prototype c.1968. Not introduced				
3707A	BB & Sweep Generator for MLA applications	1980	1983	$2k	28	$0.06M	Repackage of BB & Sweep circuits from 3710A	30			Harry Elder
3708A	Noise and Interference Test Set	1984	1999	$22k (avg)	2400	$53M	Price went from $13.5k (1985) to $32k in 1999	50	HW	HPJ July 1987	Geoff Waters
3709A/B	Constellation Analyzer	1986	1994	$11k	800	$9M	B model superceded A model in 1987	56	HW (Plus ET constellation generator)	HPJ July 1987	Boyd Williamson/David Haworth
3710A	Microwave Link Analyzer System (70 MHz IF)	1972	1985	$19k (avg)	3590	$68M	Includes 3715/16A, 3702B, 3703B/3705A	23	HW	HPJ September 1972	Reid Urquhart/Colin Appleyard
3711A	Microwave Link Analyzer System (70/140 MHz IF)	1979	1989	$28k (avg)	1096	$31M	Includes 3791B, 3712A and 3793B	25			Ivan Young
3712A	MLA Receiver (70/140 MHz IF)	1979	1989				See 3711A System				Ivan Young
3715A	MLA Baseband & Sweep Plug-in (LF Test Tones)	1972	1985				See 3710A System			HPJ September 1972	Alastair Sharp
3716A	MLA Baseband & Sweep Plug-in (LF & HF Test Tones)	1972	1985				See 3710A System		HW	HPJ September 1972	Alastair Sharp
3717A	FM Modulator/Demodulator (70 MHz IF)	1980	1991	$15k	853	$13M	Used Mod and Demod from Northern Telecom	42			David Haworth
3720A	Spectrum Display for Correlator	1972	1979	$7.3k	220	$1.6M	Price went from $6k in 1972 to $8.7k in 1979	255		HPJ November 1972	David Morrison
3721A	Correlator	1968	1979	$9.5k	610	$5.8M	Price went from $8.5k in 1968 to $11.2k in 1979	250	HW	HPJ November 1969	Brian Finnie/George Anderson
3722A	Digital Noise Generator	1967	1980	$3.3k	1590	$5.2M	Price went from $2.6k in 1968 to $4.2k in 1980	246	HW	HPJ September 1967	Gordon Roberts/Finnie/George Anderson
3724A	Baseband Analyzer System	1981	1984	$40k (avg)	103	$4.1M	Includes 3725A and 3726A	31		HPJ April 1982	Richard Roberts/Guy Douglas
3725A	Baseband Analyzer Display	1981	1984				Included in 3724A				
3726A	Filter Unit (with plug-in filters, 3726XA)	1981	1984				Included in 3724A				
3730A	Microwave Downconvertor for MLA (70 MHz IF)	1972	1980	$3.7k	994	$3.7M		27		HPJ September 1972	Mike Crabtree
3730B	Microwave Downconvertor for MLA (70/140 MHz dual-IF)	1980	1991	$3.8k	1243	$4.7M		29			James Robertson
3734A	Frequency Counter (5-digit, 2 MHz)	1963	1969	$1k			Designed in Bedford. Initially product 5534A				Don Summers
3735A	Frequency Counter (6-digit, 12.5 MHz)	1965	1969	$1.5k			Designed in Bedford				Alistair Lucas
3736A - 39A	Plug-ins for 3730A (4, 8, 11, 14 GHz)	1972	1980	$3.2k (avg)	1600	$5.1M	Volumes approx 1.6 times 3730A				Mike Crabtree
3736B - 39B	Plug-ins for 3730B (4, 8, 11, 14 GHz)	1980	1991	$7.5k (avg)	2000	$15M	Volumes approx 1.6 times 3730B				James Robertson
3740A	IF Switch (70 MHz), MLA Accessory	1973		$0.41k				45			
3743A	IF Amplifier (70MHz)	1976	1985	$1.4k			Price went from $800 in 1976 to $1.6k in 1985	45			
3744A	Baseband Sweeper Accessory	1975	1981	$1.9k	170 (est)	$0.3M	Generated BB Sweep to 15 MHz from MLA IF	45			Peter Rigby
3745A	Selective Level Measuring Set (25 MHz)	1974	1982	$21k	374	$8M		69		HPJ January 1976	Reid Urquhart
3745B	Selective Level Measuring Set (25 MHz)for BELL Std.	1975	1982	$21k	144	$3M		69	NMS T.2010.78		Reid Urquhart
3746A	Selective Level Measuring Set	1982	1991	$15k	1502	$22.5M	Price went from $20k in 1982 to $14k in 1984	81	HW		Boyd Williamson
37461A	Display	1982	1988	$3.2k	153	$0.5M	Accessory for 3746A to display FDM scans	81			Boyd Williamson
3747A	Selective Level Measuring Set (90 MHz)	1977	1985	$27k	130	$3.5M	Price went from $22k in 1977 to $33k in 1985	78			Hugh Walker
3747B	Selective Level Measuring Set (90 MHz)for BELL Std.	1977	1985	$27k	51	$1.4M	Price went from $22k in 1977 to $33k in 1985	78			Hugh Walker
3748A	Selective level Measuring Set (4 MHz)						Not introduced	80			John Coster
3750A	Push Button Attenuator 75 ohm (0 to 100 MHz)	1970	1990	$0.45k	4700	$2.1M	Price went from $200 in 1970 to $650 in 1990	45	NMS (with MLA) and HW		
3751A	Cable Group Delay Test Set for 60 MHz cable						Based on Bell spec. c.1974 Not introduced				
3754A	Access Switch 10-way (25 MHz)	1977	1991	$2.4k	2100	$5M		79	HW, NMS T.2013.63	British Patent 1 482 290	Hugh Walker
3755A	Access Switch Controller	1976	1991	$2.3k	700	$1.6M		80			David Dack/Kevin Bradford
3756A	Access Switch 10-way (90 MHz)	1978	1991	$3.4k	892	$3M		80			David Haworth
3757A	Access Switch 10-way (8.5 MHz)	1978	1991	$1.1k	1021	$1.1M		80			Hugh Walker
3760A	Pattern Generator (150 Mb/s)	1973	1980	$6.2k	313	$2M		152		HPJ November 1973	John Stinson
3761A	Error Detector (150 Mb/s)	1973	1980	$5.2k	313	$1.6M	No separate unit numbers so same as 3760	152			Tom Crawford/James Robertson
3762A	Pattern Generator (150 Mb/s)	1977	1987	$9k	1170	$10.5M	PO Tester 273 & PO Tester 249 (120 Mb/s)	156	HW (PO Tester 273A)	USPatent 4189621 P. Scott	James Robertson

A1 HPSQF Product Catalogue 1965 to 2005

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
3763A	Error Detector (150 Mb/s)	1977	1987	$9k	1100	$9.9M	PO Tester 274 & PO Tester 250 (120 Mb/s)	156	HW (PO Tester 274A)		Bob Thomson
3764A	Multirate BER Test Set (2/8/34/140 Mb/s)	1984	1996	$15k (avg)	4700	$70M	Option 007 with 140 Mb/s Jitter	160	NMS T.2010.69 (Jitter version), HW (Std.), BJ (Opt001)		Bob Thomson/David Easingwood-Wilson
3769A	General Purpose BER Tester (50 Mb/s)						Became 37800A with plug-in. Not introduced				John Coster
3770A	Amplitude/Delay Distortion Analyzer	1974	1980	$7k	429	$3.6M	US Patent 3970926, UK Patent 1429617	96		HPJ November 1974	David Guest
3770B	Telephone Line Analyzer	1976	1987	$9.5k	1190	$11.3M	PO Tester 56	102	HW		Rajni Patel (Kansagra)
3771A/B	Data Line Analyzer	1978	1986	$8.3k	522	$4.3M	Measures level, noise and transients	102			
3773A	Transmission Impairments Test Set (TIMS)						Became 4947A at launch (see below)				
3774A	In-service TIMS						Became 4948A at launch (see below)				
3776A	PCM Terminal Test Set	1983	1994	$11k	1439	$15.8M		146	HW		Rob Pearson/Andy Batham
3776B	PCM Terminal Test Set for NA std.	1983	1994	$11.3k	1465	$16.5M		146	NMS T.2010.68		Rob Pearson/Andy Batham
3777A	Channel Selector (for 3779)	1977	1994	$4.3k	1650	$7M		143			
3779A	Primary Multiplex Analyzer	1978	1981	$24k	192	$4.6M		139		HPJ January 1980	Rod MayRob Pearson
3779B	Primary Multiplex Analyzer (BELL Std.)	1978	1981	$24k	277	$6.6M		139			Rod May/Rob Pearson
3779C	Primary Multiplex Analyzer	1981	1994	$27k	1060	$28.6M		144	BJ		Mark Dykes
3779D	Primary Multiplex Analyzer (BELL Std.)	1981	1994	$27k	600	$16.2M		144			Mark Dykes
3780A	Pattern Generator and Error Detector (50 Mb/s)	1975	1991	$8.5k	4516	$38M		155	NMS T.2010.66	HPJ March 1976	Ivan Young
3781A	Pattern Generator CEPT (50 Mb/s)	1980	1990	$8k	937	$7.5M		157	HW		Aileen Appleyard
3781B	Pattern Generator BELL (50 Mb/s)	1980	1990	$7.5k	1310	$9.8M		157			Bob Thomson
3782A	Error Detector CEPT (50 Mb/s)	1980	1990	$7.2k	921	$6.6M		157			Aileen Appleyard
3782B	Error Detector BELL (50 Mb/s)	1980	1990	7k	1345	$9.4M		157			Bob Thomson
3783A	Frame Analyzer (2 Mb/s)	1980	1984	$2.2k	55	$0.12M		158			Aillen Appleyard
3784A	BER & Jitter Test Set (2, 8, 34 Mb/s)	1989	1996	$16k (avg)	2200	$35M	$13k for BER, $20k for BER + Jitter	163			Peter Scott
3785A	Jitter Test Set CEPT (2, 8, 34 Mb/s)	1981	1990	$13.5k	850	$11.4M		158	NMS T.2010.67		Peter Scott
3785B	Jitter Test Set BELL (1.5, 6.3, 44 Mb/s)	1981	1990	$15k	900	$13.5M		158			Peter Scott
3786A	Regenerator test Set (2 Mb/s)						Not introduced				Ivan Young
3787B	Digital Data Test Set	1987	1994	$9k	726	$6.5M		160	NMS T.2010.77		James Robertson
3788A	Primary Rate Tester (2 Mb/s)	1989	1993	$3k	333	$1M	Badge engineered box from Necsy (Italy)	164			
3789A/B	DS3 Test Set (Bell Std.)	1987	1994	$12k	1050	$12.5M	3789A was reduced spec. but sold only 38	160			Peter Scott
3790A	Microwave Link Analyzer System (140 MHz IF)	1974	1979	$21k	71	$1.5M	Includes 3791A, 3792A and 3793A	24			Ian Matthews
3791A	MLA Baseband & Sweep Plug-in	1974	1979				See 3790A System				
3791B	MLA Baseband & Sweep Plug-in	1979	1989				See 3711A System				
3792A	MLA Receiver (140 MHz IF)	1974	1979				See 3790A System				
3793A	Group Delay & Differential Phase Plug-in	1974	1979				See 3790A System				
3793B	Group Delay & Differential Phase Plug-in	1979	1989				See 3711A System				
3795A	Third Generation Microwave Link Analyzer						Not introduced	25			Richard Roberts
4934A	Transmission Impairments Measuring Set (TIMS)	1989	2000	$4.1k	1750/yr♣	$57M	Probably sold over 16000 units, originally 3772A	111	NMS T.2010.71		Andy Batham
4947A	Transmission Impairments Measuring Set (TIMS)	1986	1994	$8k	767	$6M	Known as "MacTIMS" originally 3773A	109			Mike Bryant?
4948A	In-service TIMS (used modem signal, V.26/27/29, and V.33)	1986	1994	$13.5k	1120	$15M	Known as "ITIMS", originally 3774A	105			Norman Carder/Gordon Rhind
5090A/B	Standard Frequency Receiver		1967				Used 200 kc/s Droitwich transmitter signal				Brian Malcolm
5263A	Time Interval Plug-in	1965	1969				Developed from 5262A plug-in to 5245L Counter				Alistair Lucas
SYSTEMS PRODUCTS											
37013A	FDM System Software (HP 1000 based)	1978	1983				Replaced by 37016A	85			Dave Warren
37014A/B/C/R	FDM System Software (PC based)	1980	1991		9		B version from 1984, various desktops	85			Peter Locke
37016A/R	FDM System Software (HP 1000 based)	1983	1991		7		Part of 37050S, ran on HP 1000 A-series	85	MH		Vince Butler
37018A/B	Baseband Analyzer System Software	1984	1984	$2.5k			NPR V-curves plotted/analyzed on HP85				Guy Douglas
37050S	FDM Network Monitoring System (HP 1000)	1984	1991					85			Vince Butler
37051S	FDM Network Monitoring System (HP Series 300)	1984	1991				Used 37014B/C software	85			
37080A/B/C	Noise and Interference Test System Software (3708S)	1984	1988	$2.5k	32	$0.08M	Controlled 3708A, A/B/C variants match BERT	50			
37480A	Openview Data Line Monitor	1989	1990		4		Controls 4948A ITIMS and 3777A switch	108		HPJ April 1990	Mike Hurst

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
37100S	**RATES System**	1983	1994				Four main systems (UK, France, S Africa and Kentucky Emergency Warning System (KEWS)	115	MH		Robert Duncan/Lawrence Lowe/Gregan Crawford/Peter Locke/Ian Burrows/Mike Hurst
37011A/R	TSS/1000 Telecom System Supervisor	1983	1994				For FDM & RATES s on HP 1000 A-series		MH		Vince Butler
37031A	Control Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Lawrence Lowe/Peter Locke/Mike Hurst
37033A	Automatic Test Point Assignment Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Mike Hurst
37034A	Circuit Record Card Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Lawrence Lowe/Peter Locke/Mike Hurst
37035A	Fault Clear System Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Lawrence Lowe/Peter Locke/Mike Hurst
37036A	ART/1000 Circuit Auto-routining Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Lawrence Lowe/Peter Locke/Mike Hurst
37037A	External Instrument Control Software	1983	1994				RATES Software runs on HP 1000 A-series		MH		Lawrence Lowe/Peter Locke/Mike Hurst
37130A	**Twin Access Module (TAM) Card Cage**	1983	1994	$3k							Gregan Crawford
37132A	TAM Control Card	1983	1994	$600					MH (in 37140 card cage)		Gregan Crawford
37133A	TAM Bus Card	1983	1994	$210							Gregan Crawford
37140A	**Access & Test Combined Unit Card Cage**	1983	1994	$2.8k					MH		Robert Duncan
37150B	**Access Module Card Cage**	1983	1994	$2.1k							Gregan Crawford
37151A	Access Card	1983	1994	$860					MH		Gregan Crawford
37151B	Access Split & Terminate	1983	1994								Gregan Crawford
37152A	Control Card	1983	1994	$1.25k					MH		Gregan Crawford
37153A	Bus Card	1983	1994	$660							Gregan Crawford
37154A	Loop-back, Oscillate & Terminate (LOT) Card	1983	1994								Gregan Crawford
37156A	Jack Access Card	1983	1994								Gregan Crawford
37160B	**Test and Measurement Unit (TMU) Card Cage**	1983	1994	$2k							Robert Duncan
37161A	TMU Control Card	1983	1994	$1.6k							Robert Duncan
37161C	TMU Control & Modem Card	1983	1994	$2k					MH		Robert Duncan
37161D	TMU Control & Modem Card	1983	1994								Robert Duncan
37162A	Input Card	1983	1994	$1.2k					MH		Robert Duncan
37162D	Input & Termination Card	1983	1994								Robert Duncan
37163A	Oscillator Card	1983	1994	$1.8k					MH		Robert Duncan
37165B	Monitor Card (North American)	1983	1994								Robert Duncan
37165C	Monitor Card (UK)	1983	1994	$1.6k					MH		Robert Duncan
37165D	Monitor & Oscillator Card	1983	1994								Robert Duncan
37166A	Level Measurement Card	1983	1994	$2.2k					MH		Robert Duncan
37166B	Enhanced Level Measurement Card	1983	1994								Robert Duncan
37166D	Level Measurement & Digital Multimeter Card	1983	1994								Robert Duncan
37167A	Termination card	1983	1994	$1k							Robert Duncan
37168A	Digital Multimeter Card	1983	1994	$1.5k					MH		Robert Duncan
37169A	Signalling Card	1983	1994	$1.9k					MH (in 37140 card cage)		Robert Duncan
37169B	Signalling Card (North American)	1983	1994								Robert Duncan
37170A	External Instrument Card	1983	1994								Robert Duncan
37170B	External Instrument & Termination Card	1983	1994	$1.1k			Affectionately known as the "exterminate" card!		MH		Robert Duncan
37171B	Signalling Card (North American)	1983	1994						MH (in 37140 card cage)		Robert Duncan
37178A	Power Supply Card	1983	1994	$1.4k			Separate power supply rack module, rather than a card		MH (x2)		Gregan Crawford
37180B	**Chain Selector Card Cage**	1983	1994	$4.8k							Gregan Crawford
37181A	Chain Select Control Card	1983	1994	$1.1k							Gregan Crawford
37190A	**System Manager Training**						Eight day course				
37191A	Operator Training Course						Three day course for 37100S users				
37192A	Measurement Subsystem Training						Three day course on installation and maintenance				
HP-IB EXTENDERS & MODEMS											
37201A	HP-IB Extender (WAN)	1979	1991	$2.4k	4200	$10M	Worked over Modem link or twisted pairs	223	MH (Opt. 50 version for FDM system)	HPJ August 1979	David Guest
37203A/L	HP-IB Extender (Coax or fibre over 1000m)	1980	1988	$1.4k	24000	$33M	Estimate 20% with Optical option at $470	226		US Patent 4,451,886	David Guest
37204A/B	HP-IB Extender (Coax or fibre over 1000m)	1988	1999	$1.6k (avg)	26500	$42M		229	NMS T.1986.614-617	US Patent 4,852,043	David Guest

A1 HPSQF Product Catalogue 1965 to 2005

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
37210T	Synchronous Modem (4.8 kb/s)	1981	1983	$4.3k	281	$1.2M	OEM Modem Card from Paradyne	232			
37220T	Synchronous Modem (9.6 kb/s)	1981	1983	$6.1k	367	$2.2M	OEM Modem Card from Paradyne	232			
37212A	Full-duplex Modem (1200 bps, V.22)	1983	1988	$1.1k	4950	$5M	Auto-dialling and auto-answer features	235			Dave Warren
37212B	Full-duplex Modem (2400 bps, V22bis, error-correcting)	1987	1991	$0.92k	9000	$8M	Auto-dialling and auto-answer features	235	MH		Eric Percival
37213A	Full-duplex Modem Card (1200 bps, V.22)	1983	1990	$1.1k	2820	$3M	Card version of 37212A for 37214A				
37214A	Systems Modem Card Cage	1983	1990	$1.7k	1670	$2.8M	8 slots, one reserved for control card				
37215A	RS232C interface card for external modem	1983	1990	$0.2k	6400	$1.2M					
37216A	Terminal interface card	1983	1990	$0.2k	1430	$0.3M					
37222A	Integral Modem (1200 bps, V.22)	1984	1986	$1.35k			Goes inside HP1000 A series computer				Dave Warren
37230A	Short Haul Modem	1981	1985	$1.2k	765	$0.9M	Operates over copper pairs to 19.2 kb/s	237			Kevin Bradford
HORNET											
37701A	T1 Tester	1990	1993	$7.5k	433	$3.2M	Replaced by 37701B	173			
37701B	T1/Datacom Tester	1993	2000	$8.5k	155/yr♣	$9.2M	Estimate 1050 units shipped	173			
37702A	Digital Data Tester	1992	2000	$12k	200/yr♣	$19M	Does T1/FT1/DDS and Datacom (Est. 1300 units)	173			
37704A	SONET Test Set (up to STS-12, 622 Mb/s)	1992	1996	$45k	75/yr♣	$16M	Estimate around 300 units shipped	175			
37711A	Digital Data Tester	1990	1993	$9k	400	$3.5M	Replaced by 37701B Opt 002 & 37702A	173	HW		
37714A	PDH/SDH/ATM Tester up to 622 Mb/s	1993	1997	$35k (est)	150	$5M	Replaced by 37717B	179			Tommy Cook
37717A	PDH/SDH/ATM Tester up to 622 Mb/s	1993	1995	$40k (est)	1378	$55M	Replaced by 37717B	179			Tommy Cook
37717B	PDH/SDH/ATM Tester up to 622 Mb/s	1995	1999	$48k*	2200	$103M	Replaced by 37717C	179			Tommy Cook
37717C	PDH/SDH/ATM Tester up to 2.488 Gb/s	1995	2000	$42k*	3509	$147M	2.488 Gb/s (STM-16) in 1998, colour display	179			Tommy Cook/Bob Thomson
37718A	PDH/SDH/ATM Tester up to 2.488 Gb/s + Jitter	1998	2004	$115k*	3800	$380M	Codename "Firefly", Jitter optional	180	NMS T.2010.72		Bob Thomson
37718B	PDH/SDH/ATM Tester up to 622 Mb/s	2000	2003	$48k*	1220	$52M	Codename "Locust"	180			Bob Thomson
37718C	PDH/SDH/ATM Tester up to 155 Mb/s	2000	2003	$42k*	1370	$57M	Lower specification versions of "Firefly"	180			Bob Thomson
37719A/B/C	SONET only version of 37718A	1999	2002	$80k*	590	$47M		181			Tommy Cook
37720A	SONET/SDH Tester up to 2.488 Gb/s	2001	2003				Low cost solution for tributary testing	181			Tommy Cook/ Malcolm Paterson
37721A	Digital Transmission Analyzer (PDH to140 Mb/s)	1990	1999	$12k	3500	$47M	704 kb/s to 140 Mb/s, BER only	172			John McElroy
37722A	Digital Telecom Analyzer (CEPT rates)	1992	1998	$8k	180/yr♣	$10.5M	64, 704, 2048, 8448 kb/s (Est.1300 units)	173	HW, Queensferry Museum		
37724A	SDH/PDH Test Set (155 & 622 Mb/s, 2 - 140 Mb/s)	1992	1998	$50k (avg)	1500	$75M	Includes 37772A and 37776A optical modules	175	HW, NMS T.2013.62		Peter Scott
37725A	SONET/SDH for Manufacturers (to 2.488 Gb/s)	2001	2003				Differential electrical interfaces. Optional Jitter	181			
37729A	Frame Generator (2, 8, 34, 140 Mb/s)	1991	1994	$22.2k	39	$0.8M	Badge-engineered box from ICT (Spain)	166			
37730A	Frame Analyzer (2, 8, 34, 140 Mb/s in-service)	1991	1994	$17.7k	150	$2.6M	Badge-engineered box from ICT (Spain)	166			
37732A	37722A Telecom Analyzer with added Datacom	1992	2000	$10.5k	630/yr♣	$50M	Estimate 5000 units shipped	173	HW		Alistair Reynolds
37741A	DS1 Tester	1991	2000	$4.1k			Handheld tester developed by Taq (contract)	165			
37742A	2M Test Set	1992	2000	$5.6k*	400/yr♣	$16M	Badge-engineered Handheld from ICT (Spain)	166			
37743A	DS3 Test Set low cost	1992	1994	$7.5k	100?		Badge-engineered box from T-Com (USA)	166			
37744A	Sonet Test Set (up to OC-3, 155 Mb/s) low cost	1992	1993	$9.9k	100?		Badge-engineered box from Antel (USA)	166			
37772A	STM-0/STM-I Optical Interface Module for 37724A 1310nm	1992	1998	14.5k	450		Included in 37724A and 37704A		HW		
37776A	STM-1/STM-4 Optical Interface Module for 37724A 1310nm	1992	1998	25.5k	1350		Included in 37724A and 37704A				
37778A	STM-16/OC-48 Test Set	1996	1999	$118k*	400	$47M	VXI Series 90 Cards repackaged E1675/1669	177			Peter Scott
E4540A	Distributed Network Analyzer Software	1994	2004	$7k (avg)	35/yr♣	$2.2M	Controlled 377XX products remotely			HPJ October 1994	Vince Butler
E4547A	MTIE/TDEV Wander Application Software	1999	2003				Worked with 37718A. Codename "Einstein"				
SIGNALING TESTERS for SS7											
37900A	Signaling Test Set						First version using HP Vectra computer	199			David Guest
37900B	Signaling Test Set	1989	1991		150		B/C used same software/interface cards	199			David Guest
37900C	Signaling Test Set	1989	1991		170		C version used R/332 computer in portable package with fewer signaling links	202			David Guest
37900D	Signaling Test Set	1991	1998	$100k	266/yr♣	$180M	Estimate 1900 total units.	204	NMS T.2010.70		Dave Warren
37915A	2 Mb/s Signaling/Interface Card	1991	1998		570/yr♣		Included in 37900D				Ian Burrows
37916A	1.5 Mb/s Signaling/Interface Card	1991	1998		180/yr♣		Included in 37900D				Ian Burrows
37918A	DS-0/V.35 Signaling/Interface Card	1991	1998		190/yr♣		Included in 37900D				Ian Burrows

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
37920A	Upgrade to Four Signaling Links	1991	1998		220/yr♣		Included in 37900D				Dave Warren
37900E 37907A	Signaling Advisor	1998		$27k*			Based on J2300 Internet Advisor (WIN 95). In 2001 transferred out of SQF to Colorado	213	HW		Ivan Young
37908A	Signaling Advisor Software	1998		$26k*			Many options for different applications				
E7571-5A	Line Interface Modules for 37900E	1998		$12k*			Developed by ECS at SQF under contract		HW		Ivan Young/David Mackenzie
E7577A	Line Interface Module Undercradle	1998									Ivan Young
HIGH-SPEED BER TESTERS											
70311A	3.3 GHz Synthesized Clock Source Module	1991	1999	$30k	729		Used with 71603A/B and 71604A/B				John Coster
70312A	1.5 GHz Synthesized Clock Source Module	1991	1995	$17.5k	6		Used with 71601A				
71601A	1 Gb/s Error Performance Analyzer	1990	1993	$64.3k	14	$0.9M					Jim Kendall
71602A	1 Gb/s Pattern Generator	1990	1993	$41.6k	12	$0.5M					Jim Kendall
71603A	3 Gb/s Error Performance Analyzer	1990	1993	$99.5k	106	$10.5M	"Daedalus" Replaced by B version in 1993	183			Jim Kendall
71603B	3 Gb/s Error Performance Analyzer	1992	2000	$115k	1035	$110M	"Theseus"	183			Jim Kendall
71604A	3 Gb/s Pattern Generator	1990	1993	$79k	9	$0.72M	Replaced by B version in 1993				Jim Kendall
71604B	3 Gb/s Pattern Generator	1993	2000	$80k	20	$1.6M					Jim Kendall
70841A	3 Gb/s Pattern Generator Module	1990	1993	$29.9	127		Included in 71603A and 71604A				Jim Kendall
70841B	3 Gb/s Pattern Generator Module	1993	2000	$44k	1180		Included in 71603B and 71604B				Jim Kendall
70842A	3 Gb/s Error Detector Module	1990	1993	$29.6k	115		Included in 71603A				Jim Kendall
70842B	3 Gb/s Error Detector Module	1993	2000	$39k	1146		Included in 71603B		HW		Jim Kendall
70845A	1 Gb/s Pattern Generator Module	1990	1993	$19.5k	27		Included in 71601A and 71602A				Jim Kendall
70846A	1 Gb/s Error Detector Module	1990	1993	$16k	15		Included in 71601A				Jim Kendall
71612A	12 Gb/s Error Performance Analyzer	1993	1998	$400k	160	$60M	Codename "Ariadne", superseded by B/C	184			Jim Kendall
71612B/C	12 Gb/s Error Performance Analyzer	1998	2003	$400k	540	$210M	Includes 70843 plus 70340 in MMS display	184	NMS T.2010.73.1-2		Jim Kendall
70843A/B/C	12 Gb/s Pattern Generator and Error Detector	1993	2003	$120k*			Included in 71612A/B/C				Jim Kendall
HP 75000 Series 90 Modular SONET/SDH Analyzer								176			
Series 90	Control Software and Integration										Dave Warren
E1615A	ATM Generator/Receiver (2488 Mb/s)	1995	2000	$199k*	4/yr♣		Bundled System (ATO)			HPJ December 1996	
E1650A/B	SONET/SDH Analyzer (to 622 Mb/s)	1991	2000		25/yr♣		Bundled System				
E1651B	System Panel (PC User Interface)	1991	2000	$22.5k*	12/yr♣		HP Vectra PC (486) with software installed				
E1652A	SONET/SDH Analyzer (52/155 Mb/s)	1991	2000		5/yr♣		Bundle				
E1654A/B	SONET/SDH Analyzer (to 2488 Mb/s)	1992	1994				Bundle				
E1654D	SONET/SDH Analyzer (to 2488 Mb/s)	1994	2000		5/yr♣		Bundle				
E1654E	SONET/SDH Analyzer (to 2488 Mb/s)	1994	2000		7/yr♣		Bundle				
E1654F	SONET/SDH Analyzer (to 2488 Mb/s)	1995	2000		10/yr♣		Bundle				
E1655A	ATM Analyzer (155 Mb/s)	1992	2000				Bundled System from ATO with some QTD modules				
E1659A	ATM Analyzer Software	1991	2000								
E1661A	Optical Interface module (52/155 Mb/s)	1991	2000		38/yr♣						
E1662A/B/C	Optical Interface Module (155/622 Mb/s)	1991	2000	$27.6k*	14/yr♣						
E1663A	Electrical Interface Module (52/155 Mb/s)	1991	2000	$6.5k*	90/yr♣		Transceiver doing clock recovery and code				
E1665A	Quad O/E E/O Converter (155 Mb/s)	1991	2000	$42.7k*							
E1666A/Z	EOC Access Module	1993	2000	$5.7k*	12/yr♣						
E1667A	Jitter Analyzer (52/155/622 Mb/s)	1993	2000	$53k*	55/yr♣		Generates/measures jitter/wander. Designed by Microwave Logic				Ivan Young
E1668A	Jitter Analyzer (2488 Mb/s)	1994	2000	$37k*	20/yr♣						
E1669A/B	Optical Interface Module (52/155/622/2488 Mb/s)	1994	2000	$60k*	40/yr♣		1550 or 1300 nm Transceiver for E1675/67/68				
E1670A	Optical Interface Module	1994	2000	$41k*	10/yr♣						
E1671A	Transport Overhead (TOH) Generator (52/155/622 Mb/s)	1991	2000	$31.5k*	8/yr♣						ATO then Malcolm Paterson
E1672A	Transport Overhead Receiver (52/155/622 Mb/s)	1991	2000	$31.5k*	8/yr♣						ATO then Malcolm Paterson
E1673A	Transport Overhead Generator (52/155 Mb/s)	1991	2000								
E1674A	Transport Overhead Receiver (52/155 Mb/s)	1991	2000								

A1 HPSQF Product Catalogue 1965 to 2005

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
E1675A	Transport Overhead Analyzer (2488 Mb/s)	1992	2000		40/yr♣		Transceiver with Tributary Add/Drop capability				Bob Thomson/Dave Warren
E1676A/B	Transport Overhead Analyzer (52/155/622/2488 Mb/s)	1994	2000	$50k*	15/yr♣						Bob Thomson
E1679A	Timing Reference Module	1991	2000	$6.8k	110/yr♣		Reference clock from a variety of inputs				Ivan Young
E1681A	Synchronous Payload Envelope (SPE) Generator	1992	2000		5/yr♣						ATO then Malcolm Paterson
E1682A	Synchronous Payload Envelope (SPE) Receiver	1992	2000		5/yr♣						ATO then Malcolm Paterson
E1683A	Payload VT/TU Generator	1991	2000	$5.5k*	80/yr♣						James Robertson/Malcolm Paterson
E1684A	Payload VT/TU Receiver	1991	2000	$5.5k*	80/yr♣						James Robertson/Malcolm Paterson
E1691A	ATM Generator (155 Mb/s)	1992	2000		5/yr♣		Designed at ATO				
E1692A	ATM Receiver (155 Mb/s)	1992	2000		5/yr♣		Designed at ATO				
E1693A	ATM Generator (155/622 Mb/s)	1993	2000		8/yr♣						
E1694A	ATM Receiver (155/622 Mb/s)	1993	2000		8/yr♣						
E1695A	DS3 Line Interface	1992	2000		1/yr♣						
E1696A	ATM Optical Load Generator (155 Mb/s)	1993	2000		1/yr♣						
HP 75000 Series 95 SONET/SDH Tributary Test System (TS-2000)								177			
Series 95	Control Software and Integration										Dave Warren
E1550A	32-Channel Balanced Daisy Chain Switch	1994	1998	$10k*	20/yr♣						
E1551A	STS-1/DS3 Daisy Chain Switch	1994	1998	$11k*	30/yr♣						
E1552A	STS-1/DS3 Test Access Switch	1994	1998		15/yr♣						
E1553A	Dual STS-1/DS3 Amplifier/LBO	1994	1998	$7.2k*	7/yr♣						
E1554A	8-Channel 75 Ohm Daisy Chain Switch	1994	1998	$14k*	30/yr♣						
E1620A	16-way fanout module (34 Mb/s)	1996	1998		1/yr♣						
E1623A	TS-2000 SDH & ATM functional test system	1996	1998		2		Bundled System				Dave Warren
E1625A	TS-2000 SONET functional test system	1995	1998		10		Bundled System				Dave Warren
E1660A/B	32-way fanout module (DS1/E1)	1994	1998		10/yr♣						
E1664A	16-way Fanout (STS-1/DS3)	1993	1998		15/yr♣						
E1664B	8-way Fanout (STM-1/139Mb/s)	1993	1998		7/yr♣						
E1685A	STS1/DS3 Quad Transceiver	1993	1998	$23k*	30/yr♣						
E1685B	STM-1/STS-3(c)/ 139Mb/s quad transceiver	1993	1998	$30k	5/yr♣						
E1686A	DS1/DS3 Transceiver	1994	1998	$15.5k*	113						
E1686B	2/8/34/139 Mb/s Transceiver	1994	1998	$21k*	6/yr♣						
E4502A	Single 1x17 optical switch	1994	1998	$20k*	56						
E4503A	Dual 1x8 optical switch	1994	1998	$19k*	78						
E4504A	Dual 1x4 optical switch	1994	1998	$18k*	129						
SpectralBER								177			Douglas Butler
J1420A/B	10 Gb/s DWDM Receiver	2000	2001		200		Badge-engineered Ando module				
J1421A	10 Gb/s Clock Source	2000	2001		85		Badge-engineered Ando module				
J1422A/B	10 Gb/s DWDM Receiver	2000	2001		110		Badge-engineered Ando module				
J1425A	10 Gb/s SpectralBER System	2000	2001	$305k+	117	$40M+	Bundled System (sold 90% in USA)				
J1426A	2 Mb/s Clock Module	2000	2001		26		Badge-engineered Ando module				
J1427A	1.5 Mb/s Clock Module	2000	2001		4		Badge-engineered Ando module				
J4221A	2.5 Gb/s Single-rate SpectralBER System	1999	2001	$135k+	37	$5M+	Bundled System (50% Europe, 50% USA)				
J4222A	2.5 Gb/s Multi-rate SpectralBER System	1999	2001	$175k+	50	$9M+	Bundled System (50% Europe, 50% USA)				
J4225A	2.5 Gb/s Short Reach DWDM Receiver Module	1999	2000								
J4226A	2.5 Gb/s Long Reach DWDM Receiver Module	1999	2001								
J4227A	2.5 Gb/s Multi-rate DWDM Receiver Module	1999	2001				Multi-rate 155/622/2488 Mb/s				
J4230A	2.5 Gb/s Single-rate 1310nm DWDM Transmit Module	1999	2000								
J4231A	2.5 Gb/s Single-rate 1550nm DWDM Transmit Module	1999	2000								
J4232A	2.5 Gb/s Single-rate 1550nm ITU-T DWDM Tx Module	1999	2000								
J4233A	2.5 Gb/s Multi-rate 1310nm DWDM Tx Module	1999	2001				Multi-rate 155/622/2488 Mb/s				
J4234A	2.5 Gb/s Multi-rate 1550nm DWDM Tx Module	1999	2001				Multi-rate 155/622/2488 Mb/s				
J4235A	2.5 Gb/s Multi-rate 1550nm ITU-T DWDM Tx Module	1999	2001				Multi-rate 155/622/2488 Mb/s				

Product Number	Product Description	Intro	Discont'd	Approx. Price ($)	Units Sold	Volume Gross ($)	Comments	Vol. 2 Page No.	Unit in Preservation	HP Journal Article Patent	Project Leader
FINAL TRANSMISSION TEST PRODUCTS											
E7580A	ProBER 2, 2 Mb/s Handheld	1999	2003	$4.75k*	4000	$18M	Developed under contract by ECS at SQF	187	NMS T.2010.76		David Mackenzie
J2126A	Transmission Test Set (PDH, SONET/SDH to 2.5 Gb/s)	2001	2004	$53k (avg.)	1065	$57M	Codename "Lynx"	187			Malcolm Paterson
J2127A	Transmission Test Set (As for J2126A plus GbE & 10 Gb/s)	2002	2004				Codename "Cobra" (Unit numbers in Lynx)	187	HW		Jim Elliott
J7230A/B	OmniBER OTN 10G (PDH, SONET/SDH to 10 Gb/s)	2002	2005	$148k	100	$15M	Codename "Panther"	189	NMS T.2010.75		Norman Carder
J7230C	OmniBER OTN 40G (PDH, SONET/SDH up to 40 Gb/s)			$350k			Not introduced. Codename "Quantum"	190			P J Hillery/John Coster
J7231B	OmniBER OTN 10G (BER & Jitter to 10 Gb/s)	2002	2005	$146k	170	$25M	Codename "Merlin"	190			Eileen Christie
J7232A	OmniBER OTN 2.5G (BER & Jitter to 2,5 Gb/s)	2001	2005	$95k	170	$16M		190			
J7241/42A	OmniBER XM (10 Gb/s)	2002	2004	$57k	80	$4.5M	Multi-port modular. Codename "Yellowstone"	191			Hermine Schnetler/ David Newton
J7244/45A	OmniBER XM (2.5 Gb/s)	2002	2004	$48k	150	$7.2M	Codename "Nugget"	191			Hermine Schnetler/ David Newton

NOTES

Data Sources: Product Support Obsoletion Notices, Sales Stats, HP Catalogs, Facility Reviews, Contemporary Marketing Reports

Unit Volumes: As recorded in documents, e.g. Osoletion Notices. ♣ Where XX/yr shown this is the average for the 5 years FY92 to FY96 calculated accordingly for years manufactured

Where no specific total is available, I estimate total volume based on average units/yr over the lifetime of the product with a taper applied for the first and last years of 50% of the average. Gross $ Volume is number of units multiplied by approximate US List Price (see below).

Prices: USA List Price taken from HP Catalog at roughly mid-life taking some account of options where known.
Prices marked * are taken from a February 1999 spreadsheet showing the average price used by accounts
Prices for 37100S modules taken from list prices in 1989 BT RATES Tender
Prices for post 2000 products taken from contemporary marketing information
Prices marked (avg) is a ballpark average for the option mix where the options are expensive, or for products with a large price variation over the life of the product e.g. 3710 system or 3708A

Preservation: NMS - National Museums of Scotland Collection, Granton, Edinburgh
HW - Hugh Walker

MH - Mike Hurst
BJ - Bernd Johann, Germany: Museum fur Nachrichtentecknik www.museum-nt.de/

Compiled by Hugh Walker, November 2010. Updated March 2021

Appendix 2

Bibliography and HP Journal Articles

1. Bibliography

This is a short list of books on telecommunications tying in with the product families, and the times and technologies of the South Queensferry instruments.

"Communications Network Test & Measurement Handbook" by Coombs & Coombs, McGraw-Hill 1998, ISBN 0-07-012617-8

A huge, comprehensive tome of over 800 pages, compiled by Clyde Coombs (a former HP employee) and his daughter. All chapters are written by HP engineers, many from South Queensferry. A very useful book as it reflects the HP perspective on network technology and measurements in the late 1990s, shortly before the Queensferry site reached its zenith. Frequently referenced in the later chapters.

"Introduction to Telephones and Telephone Systems" by Michael Noll, Artech House 1999, ISBN 1-58053-000-1

First published in 1985 with a third edition in 1999, this book provides useful background on telecom systems from a North American viewpoint. It looks at technology evolution and the history of the North American Bell System with perspectives through to the late 1990s. Around 370 pages.

"Telecommunication Networks" by J.E. Flood, Institution of Electrical Engineers, 1997, ISBN 0-85296-886-8

This is a very useful book looking at all aspects of the network including services, switching, signalling and transmission from the European/UK perspective. First published in 1975 it was updated extensively in 1997. Authored by Flood and many other contributors from British industry and universities. Part of the IEE Telecommunications series totalling 36 volumes. Over 500 pages.

**"Telecommunications Transmission Handbook" by Roger Freeman,
John Wiley 1991, ISBN 0-471-51816-6**

A popular book at the factory, first published in 1975. The one listed here is the third edition. Gives a very good survey of FDM, TDM and Microwave Radio, all very pertinent to Queensferry's business. Over 1000 pages. A fourth edition was published in 1998.

**"Telecommunications Measurements, Analysis and Instrumentation" by Dr. Kamilo Feher and Engineers of Hewlett-Packard,
Prentice-Hall 1987, ISBN 0-13-902404-2**

This is a valuable record of technical papers written mainly by South Queensferry engineers in the 1980s, with editing and contributions from Dr. Feher. Much of the material started life as papers written for the touring telecom symposiums organised by the Queensferry marketing department in the 1980s, which travelled all over the world particularly in Europe and the USA. Referenced a number of times in later chapters, the book has over 400 pages.

**"Signalling in Telecommunications Networks" by Samuel Welch,
Institution of Electrical Engineers 1979, ISBN 0 906048 04 4**

Sammy Welch was a friend of HP South Queensferry and visited on several occasions to provide tutorials and consulting on telecommunications and especially signalling, having previously been the head of signalling at the British Post Office. He was also an expert on No. 7 Signalling, however this book does predate most common channel signalling. It gives a very good overview of the history and importance of signalling in the network.

**"Signaling System #7" by Travis Russell,
McGraw-Hill 1995, ISBN 0-07-054991-5**

A very informative book on the subject written by an engineer from one of South Queensferry's competitors, Tekelec Inc. 470 pages.

2. HP Journal Articles

Articles by South Queensferry engineers published in the HP Journal are a valuable source of information on how the products were designed. HP Journal issues are all available for download on-line as pdf files from two sources.

They are available from HP's own website at:
http://www.hpl.hp.com/hpjournal/pdfs/IssuePDFs/hpjindex.html

and also from the private HP Archive website at:
http://www.hparchive.com/hp_journals.htm

The HP Archive source appears to have slightly higher-quality scans, virtually as good as the printed copy, although file sizes are larger. For this reason, I've made most of the Chapter references using the HP Archive listing. This is complete listing of articles known to be published by the Division.

3722A Noise Generator September 1967
http://www.hparchive.com/Journals/HPJ-1967-09.pdf

3721A Correlator November 1969
http://www.hparchive.com/Journals/HPJ-1969-11.pdf

3720A Fourier Display November 1972
http://www.hparchive.com/Journals/HPJ-1972-11.pdf

3710A/02B MLA and 3730A Down Converter September 1972
http://www.hparchive.com/Journals/HPJ-1972-09.pdf

3760A/61A BERT November 1973
http://www.hparchive.com/Journals/HPJ-1973-11.pdf

3770A Telephone Line Analyzer November 1974
http://www.hparchive.com/Journals/HPJ-1974-11.pdf

3790A 140MHz MLA November 1975
http://www.hparchive.com/Journals/HPJ-1975-11.pdf

3745A SLMS January 1976
http://www.hparchive.com/Journals/HPJ-1976-01.pdf

3780A BERT March 1976
http://www.hparchive.com/Journals/HPJ-1976-03.pdf

3754A/55A Access Switch August 1978
http://www.hparchive.com/Journals/HPJ-1978-08.pdf

37201A Bus Extender August 1979
http://www.hparchive.com/Journals/HPJ-1979-08.pdf

3779A/B PCM Tester January 1980
http://www.hparchive.com/Journals/HPJ-1980-01.pdf

3724/25/26A Baseband Analyzer April 1982
http://www.hparchive.com/Journals/HPJ-1982-04.pdf

3708A C/N Test Set and 3709A Constellation Display July 1987
http://www.hparchive.com/Journals/HPJ-1987-07.pdf

4948A In-Service TIMS October 1987
http://www.hparchive.com/Journals/HPJ-1987-10.pdf

OpenView Data Line Monitor April 1990
http://www.hparchive.com/Journals/HPJ-1990-04.pdf

Virtual Remote Software October 1994
http://www.hparchive.com/Journals/HPJ-1994-10.pdf

Appendix 3

Memories of Working on the Microwave Link Analyzer Production Line in the early 1970s

by Bill Lauchlan

Bill Lauchlan joined HP South Queensferry as a test engineer in February 1967. Like a number of the early employees, he came to HP from Ferranti in Edinburgh where he had worked on the testing of radar systems for the new generation of fighter jets in the 1950s and 60s. In the early 1970s he worked on the MLA production line, and moved to Colorado Springs a few years later to provide support for Queensferry's products in North America. Later he joined the sales organization in Canada for the rest of his HP career. Bill's recollections demonstrate the valuable interaction and teamwork between design and manufacturing when the two are co-located.

I don't remember the exact year I moved into the MLA department, but it must have been around 1970. At that time, the 3701A/2A/3A was the only MLA product in production. The key people I worked with were:

Ken Nicol – Test Leader
Bill Morrison - Test Engineer
Ian Wilson – Test Engineer
Phillip Melville – Test Engineer
Ken McDougall – Production Engineer

Reliability Problems and Test Selection

Reliability was a big problem on the early MLAs. All the wires connecting to PC Boards were through push on connectors. Unfortunately they would loosen off after the instrument got bumped around a bit. Reed relays were used all over the place, and had a bad failure rate.

We used to do many overhauls on customer units to weed out these problems, by soldering all the wires onto pins, and replacing the reed relays with a newer style and test selection based on contact resistance.

We also used to spend lots of time minimizing "spurs" (spurious responses on the swept display), adjusting group delay and return loss. The IF sweep generator in the 3701A was open loop (it had no feedback control of the frequency) so we had to do lots of test selection on varactor diodes to get reasonably linear sweep across the IF band. The IF level flatness also required component test selection.

Ferrite beads seemed to be the answer to many odd problems of very high-frequency spurious oscillation, usually added to the leg of a transistor. Occasionally we would get bad batches of ferrite (probably badly fired) where the beads themselves would cause "microphonic" problems.

The original MLA Production Line used some non-HP test equipment, some of which I was familiar with from my time at Ferranti Ltd. These included the Rhode and Schwarz PolySkop (an early type of swept network analyzer with a large TV screen) and an Attenuator Test Set (I believe it was from Ad-Yu Electronics). While these were probably the state of the art at the time, they became problematic from a maintenance and calibration standpoint.

Also, we used the HP 8551A first generation Spectrum Analyzer for checking harmonics on the IF output. Again it was state of the art for the time, but most of our engineers didn't know how to use it properly.

Bell (AT&T/Western Electric) KS20548 L1/L2/L3 3701Z/2Z/3Z

While this System was the same basic platform as the first MLA, there were major mechanical and electrical/electronic improvements. Chassis metalwork was beefed-up to meet WECO environmental requirements and overall it had to be built according to WECO manufacturing Specs.

In this upgraded MLA, we got rid of many deficiencies from the original like reed relays which were replaced by solid-state diode and transistor switching, which improved reliability. Also, all coaxial connections to P.C. boards were made with gold-plated 75 ohm "Sealectro" connectors.

One interesting little problem I recall was with splitting the Baseband and Sweep Signals going into the BB + Sweep Inputs on the Receiver and 124 ohm Converters. This Circuit included a series inductor as part of a low-pass filter. When we cranked up the Sweep component to 5 volts, the Baseband Display would bow into a semi-circle. The reason for this was saturation of the ferrite core of the inductor, which I believe was about 2.2 mH. The amusing part about this problem was finding the solution. The production engineer went about sourcing a replacement inductor and brought many samples for Bill Morrison and me to test. None of them worked. In fact we stopped trying them in the equipment. We would just try picking them up with a magnet on the bench. If they lifted we would reject them. The solution was to go to an "air-cored" inductor, which we made in-house. We asked someone in the Coil/Transformer Shop to put as many turns as he could on a 1M ohm resistor. It worked like a charm and that coil carried on in the MLA design through all the new models until production stopped!

Since the Bell contract equipment was all for use in the USA, it would all be operating at 115V and 60Hz Line Frequency. Also the Sweep Frequency was 100 Hz. At that time we tested everything at 240V and 50Hz in Queensferry. 115V Operation only got a Basic Test to ensure it was wired properly and had necessary regulation . We never checked at 60 Hz. Also with a line frequency of 50Hz we got a beat with the 100Hz sweep frequency. It was decided to install 60Hz 115V on the MLA production Line. The first attempt at this used a basic construction-type generator which had a terrible waveform and we had to get a second generator with much greater capacity and a decent sine wave.

Overall, for quality, reliability and testability, this MLA design was a major improvement over the earlier version. Much of the improvement was back-engineered into the 3701A/2A/3B system.

The Second Generation 3710A/15A/16A/02B/03B/05A MLA

This MLA, while having the same core elements as the previous models, introduced Differential Phase and Gain measurement using higher Baseband Frequencies up to 8.2MHz, and with Demodulation up to 5.6MHz.

This was also the first MLA to use analogue and digital integrated circuitry. Sweep circuits and baseband circuits used digital techniques for signal generation. This allowed for lots of flexibility with baseband and sweep frequencies and also guaranteed accuracy.

I.F. frequency stability and sweep linearity were much improved in the 3710A Transmitter by using a pulse-counting discriminator in a closed loop, along with the digital sweep generation. Corresponding changes were made to the 3702B Receiver with a new X-Axis circuitry incorporating automatic phase adjustment.

The 3702B didn't benefit much from digital circuitry introduction, although there was extensive change to the demodulator loop, changing the I.F. to 17.4 MHz from 12 MHz on the earlier design and having a broader discriminator to handle higher BB frequencies up to 5.6 MHz.

Unlike the 3710A, there were no major technique changes in the 3702B, so you might expect it to be much easier to introduce to production than the 3710A. Not so, as I will elaborate on later.

The 3705A Differential Phase Detector was basically a 3703B with some added crystals to handle the higher baseband frequencies.

The Transmitter side (3710A/15A/16A) had extensive metalwork changes, moving away from a solid one piece RF Housing to a multi-piece "tin box" and moving to plug-in units for Baseband Generation (3715A/16A). Another major change was using slide-rule displays for 3710A controls. This used gears and sprockets to drive a piece of "Film Strip" to provide the read out.

Transition of 3710A System from Lab to Production

While this obviously involved both Test Engineering and Assembly/Wiring folks, I will restrict my words here to the Test Engineering side, since I didn't have so much involvement with the other parts.

The production folks for the transition were Bill Morrison and myself for the Transmitter (3710A/15A/16A) and for the Receiver side Phil Melville and Ian Harrison.

From R&D, I recall Hugh Walker, Reid Urquhart, Alistair Sharp, Bryan Lewis, Colin Appleyard, Norman Edgar, Ian Matthews, Owen Livingstone, Harry Elder, Duncan Reid, Mike Crabtree.

3710A/15A/16A IF/BB Generator

3710A - Electrical

The Main functions of Sweep and IF generation worked well. The new digitally generated sweep signal had much improved harmonic content. I don't recall any major issues in this area. The IF Generation similarly worked well with improved levelling and frequency accuracy. There were some temperature related issues with resistors being run too close to their power rating. We also had to change front panel potentiometers to high-linearity types now we had the slide-rule type scale. We did have a problem with miniature trim pots. One type had very poor setability and wouldn't retain its setting. These were eliminated.

Another component which was widely used for decoupling and inter-stage coupling was a small 2.2 µF electrolytic capacitor. This seemed to be OK in IF circuits, but was problematical in BB circuits. It was "lossy", and made the performance of otherwise stable circuits unpredictable. We had these replaced with a non-electrolytic component of the same value. The reason I remember this was that it caused some "wailing and gnashing of teeth" in either production engineering or R&D 'cos these components were much more expensive than the originals. From what I recall, manufacturing overhead was allocated on labour cost, such that saving a buck in labour could justify much more in purchased part cost. Also, the impact in warranty cost was a big lever. So the changes went ahead.

With the introduction of digital and analogue ICs, the power supplies in the 3710A were more complex than in previous MLAs. There were some teething troubles with over-voltage protection SCR's (Silicon Controlled Rectifier crowbars) not being heavily enough rated which would cause failure of the over-voltage protection after two successive fuse replacements. When this happened it could wipe out a whole bunch of chips, predominantly in the 3716A. The other power supply design problem was ground wiring. This first showed up, as I recall, as unwanted line (mains) frequency FM on the 70MHz IF signal. On investigation, the cause of this was twofold. One cause was the use of too lightweight wiring for power supply board grounds and the other was multiple ground points on the chassis. We fixed this by establishing one common ground point on the chassis and using heavy tinned-copper wiring for all the power supply and associated large capacitor grounding.

3710A – Mechanical

In early production the Gear Deck and Plug-In Fitting were two nightmare areas.

The Gear Deck was made from folded metal and bushings were aluminium. The shafting was steel of some form. This assembly should have been a precision engineered system but it was not. This gave rise to assembly folks having to custom fit all these assemblies.

Once they were installed in the main-frame, the situation was further exacerbated by the front of the 3710A not being a rigid structure. One could grab the two handles and flex the whole front of the unit. This would cause Gear Deck misalignment and also cause major Plug-In fit problems. This trouble was never addressed until development of the 140 MHz 3790 System, for which a new method was developed using precision castings, a belt drive and a bracing plate to stiffen up the front of the 3790 chassis. Technically this was successful, but for other reasons was not back-engineered into the 3710 System. It was, in my opinion, a much better arrangement.

3715A/16A – Electrical

I won’t say much about the 3715A, other than it was the low end BB Generator with only the old bottom three frequencies of the older MLA. It worked fine, but wasn’t in great demand.

The 3716A was the major Baseband Generator and used lots of new digital and phase lock loop circuitry to generate the Baseband Signal. Also it had sweep reduction circuitry to contain the FM sidebands within the Swept Bandwidth being used.

This unit worked well through introduction, with really only one significant design flaw. That was with the variable radix divider feedback loop. Due probably to manufacturing variances in some of the “digital chippery”, a race condition existed in the reset feedback loop, whereby the divider reset pulse could occur a little early and not reset properly. After trying different chips, we found that one logic gate chip had an unused element in it. So we patched the feedback loop through this unused gate, thereby introducing a little more delay into the feedback loop. This guaranteed the right guy always won the race and from then on everything worked good!

3702B I.F. / B.B. Receiver and 3705A Differential Phase/ Group Delay Detector Transition to Production

The 3705A Plug-In had no significant problems through introduction since it was basically a 3703B with some add-ons to handle the higher baseband frequencies.
Nothing more really needs saying about it.

Now the 3702B, that's another story. The primary problem areas at introduction were the demodulator AFC loop, high-voltage power supply and the X-axis deflection system.

Demodulator Loop – this consisted of three boards which were the bandpass filter and mixer to down-convert the IF signal, the local oscillator board and the discriminator board. Problems associated with this Loop were a mix of things found at introduction time and also a performance issue with AM to PM conversion which became major after introduction.

I recall the main problem was AFC capture whereby the tunable local oscillator, running open-loop prior to input signal application, would be far enough away from nominal that the resulting input to the discriminator at signal application was out of range so that the loop couldn't lock in. This required some circuit re-design to hold the LO idle frequency within range of the discriminator. I also recall some problems relating to the bandpass filter and the limiter in front of the discriminator. These are a bit foggy in my mind and were probably the beginning of seeing the AM to PM problem.

On a positive note, similar to the sweep improvements on the 3710A, the 3702B had filtering in the sweep recovery circuits to clean up the recovered sweep to be applied to the X-Axis. This alleviated an old test selection problem from the previous generation of having to test-select varactor diodes to get decent sweep linearity.

High Voltage Power Supply – with this there was also a mix of problems found at introduction and some which showed up only after customers were using the equipment in the field (particularly in Canada).

The high-voltage supply was very similar to the previous generation, however the printed circuit board layout was different. There were some tracks which were too close together. This didn't show up until customers (I believe in Alberta and the Rockies) started to have failure due to flashover at high altitude. This required board re-layout. Also, there was a ground loop problem due to PC board tracks which caused radiation from the high voltage circuit interfering with BB frequencies. Ground wiring was changed.

The most dramatic high voltage problem was where the multi-vibrator circuit would not start naturally and instead hung-up with both transistors conducting. This would cause board burn-up. I recall Mike Crabtree (R&D) got involved and we put together a little board with a low energy multi-vibrator to "kick-start" circuit and this solved the problem.

X-Deflection System – the actual deflection part of this board worked OK, but the problem area was the automatic phase-shifting circuitry. This circuitry, as I recall, used electro-optical devices (composed of an electric bulb shining on a photo-sensitive resistor). I don't now remember the specifics, but they never worked reliably.

Miscellaneous

There was a component rating problem which mostly showed on the boards in the RF Housing where the temperature was highest. We used many resistors which I believe were a metal-oxide film type. These had a characteristic where if you ran them too hot they would change value permanently, and also they could burn up. This was particularly troublesome in I.F. amplifiers where, after some time in the field, the amplifier gain would change and cause failure. I recall having a whole bunch of components up-rated after having Ohms Law checks done on the boards.

A measure which gives an idea of how well the 3702B transitioned to production was the number of production changes on the first production run – something like 150 to 250 as I recall!

A.M to P.M Conversion Problem

I don't profess to ever having completely understood this problem. I know that design changes were made in the areas of the bandpass filter and mixer and also on the Discriminator board in the limiter area. The people I associate with this were Ian Matthews and Ken McDougall. I recall that after redesign, there was a new adjustment that had to be made on the discriminator board, while a special network was inserted in the IF path between the 3710A and the 3702B. I think the network was introducing a controlled amount of AM to the IF Signal and had a known group delay response. We then had to make an adjustment on the discriminator board for a particular group delay signature, which I believe meant we had zeroed-out the AM to PM conversion in front of the discriminator. In prior testing, we were always feeding the 3702B with zero AM and therefore could not detect AM to PM conversion. (More information on this was published in the November 1975 HP Journal article about the 3790A MLA by Ian Matthews)

3710A System Retrofit Programme

Because of the reliability problems and some performance issues like AM to PM conversion, it was decided during 1973 to carry out a major retro-fit programme for early-adopter customers around the world. One of the major early adopters was the Trans-Canada Telephone System (an association of the major Telephone companies in Canada). We sold somewhere around 50 to 60 Systems (Option H50) from early production. The biggest commitment from Canada was Alberta Government Telephones who purchased in excess of 40 systems. Interestingly, Bell Canada (the biggest telephone company) only bought one to my knowledge, possibly two.

This programme involved retrofitting new high voltage boards, new demodulator loops and new X-deflection boards to the 3702B. The 3710A also had new power supply boards and a new IF attenuator installed. Various other boards were replaced 'cos of bad "trimpot" problems. Ken MacDougall carried out this programme in Europe and I was asked to implement it for North America. I spent a month (November 1973) in Canada, one week Montreal, two weeks Edmonton and one week Vancouver, carrying out this work. The retrofit programme took care of the major reliability issues, although I believe some more work had to be done to really solve the AM to PM problem.

The MLA which Bell Canada (North Bay) bought had an unfortunate history during its first six months or so. It seemed to be in our service shop more than it was in use. This was becoming a serious problem with Bell, who had the potential to buy many more systems. I asked to have this system sent to Edmonton so I could have a look at it.

On investigation, I found a major manufacturing quality problem which was the cause of the trouble. For reasons unknown to me, the assembly line had stopped wrapping wire connections on the 50 pin edge connectors holding the P.C. boards on the 3710A. Instead the wires were just being popped through the solder lugs and soldered. This gave rise to cold solder joints. On the defective Bell unit, I pulled out around a dozen wires very easily. I decided to scrap this MLA and give the customer a new one. The old unit went back to the factory and was used on the production line. When I returned to Scotland, some changes were made to prevent this problem happening again.

Manufacturing Quality and Workmanship Standards

The Western Electric Contract for the 3701Z/2Z/3Z MLA System required that these were produced in accordance with WECO Quality and Workmanship Standards. These were similar to, but in some cases more stringent than, HP comparable standards. This not only applied to wiring, but complete environmental requirements including vibration and temperature.

The MLA production line was producing equipment to both the WECO standard (for the contract equipment) and the HP standard for other commercial equipment. I believe this may have been largely responsible for concerns raised by WECO on the contract. This resulted in a visit to Queensferry by Ernie Major from the Q.A. department of HP Customer Service Centre in Mountain View, California. Ernie worked with our Q.A. department and the Production Line to bring the contract MLA production line into conformity during 1973.

On return from Canada end of 1973, I undertook a project to deal with the quality and workmanship issues on the MLA production line. I felt that having one production line working to two standards probably wasn’t the best way to be. That was changed so that everything was produced to the higher standard, which generally meant WECO.

I recall asking one of our Assembly Supervisors, who had lots of experience building to military standards, to build me an MLA (3710 System) as though it was Radar equipment. That included wrapping all wire connections. She did that very well and thereafter the Assembly Line worked to that one standard.

Production Line Test Equipment

During the evolution of the MLA products, HP also produced some new products which could improve our testing and adjustment methods. Some of these were the 140T/141T/855X Spectrum Analyzer family with Tracking Generators and also the 8601A Sweeper. I had some exposure to these through R&D using them but we didn't have any on the production line.

At some point, probably mid 1973, we had a production capacity problem. I think Sales must have sold too many (shame on them!). This required an increase in capacity on the production line. I recall having to equip something like an extra four test stations. Through the budgeting for this, I was able to find enough money to acquire some of the new Spectrum Analyzer and Sweeper equipment which were used on the MLA Line for many years thereafter. This also allowed for elimination of the old Ad-Yu Attenuator Test Set, and retiring some Polyskop equipment.

Conclusions

My involvement with the MLA Production Line finished in the first half of 1974, when John Groat left HP for the Oil Industry back in Shetland and I took over from him as Line Manager on Transferred (Licensee) products.

That didn't last long 'cos Peter Carmichael took over as Manufacturing Manager and decided I should be back in Communications. I was asked to take Line Manager, Communication Products (Analogue and Digital). Jim Stewart was Line Manager MLA products.

I think it is important to note that this account of MLA Production is my recollection some 40 years later. I was only one member of a team responsible for the manufacture of the various MLA products. I consider it a privilege to have worked along with the many dedicated people who brought their own skills and talents to bear on the various challenges presented to us.

Bill Lauchlan – August 2012

Appendix 4

Preserved South Queensferry Products

This is a listing of South Queensferry products known to be preserved as of 2013.

The main preserved group of 19 items is with the National Museum of Scotland at their Collection Centre at Granton Road, Edinburgh:
http://www.nms.ac.uk/our_museums/museums_collection_centre.aspx

The products can be viewed there by appointment. When donated in 2010, all were complete and working, and were supplied with cables and service manuals where available.

The other listings are for private collections.

National Museum of Scotland

Instrument	NMS Accession Number
3701A MLA System (3701A, 3702A and 3703B)	T.2010.65.1-3
3780A Pattern Generator and Error Detector	T.2010.66
3745B Selective Level Measuring Set	T.2010.78
3764A Digital Transmission Analyzer (Opt.007)	T.2010.69
3776B PCM Terminal Test Set	T.2010.68
3785A Jitter Test Set	T.2010.67
3787B Digital Data Test Set	T.2010.77
4934A Transmission Impairments Measuring Set	T.2010.71
37900D Signaling Test Set	T.2010.70
71612B 12 Gb/s Error Performance Analyzer	T.2010.73.1-2
37204A HP-IB Extender	T.1986.614-617
37718A OmniBER PDH/SDH/ATM Tester	T.2010.72
E7580A ProBER 2 2M handheld Test Set	T.2010.76
J7230A OmniBER OTN 10 Gb/s Test Set	T.2010.75
E5515C Wireless Communications Test Set (QMD)	T.2010.74
37724A SDH Test Set with 37772A Optical Plug-in	T.2013.62
3754A 25 MHz Access Switch	T.2013.63

3722A Noise Generator	T.2021.32.1
3721A Correlator	T.2021.32.2

South Queensferry Museum

37722A Digital Telecom Analyzer

Private Collections

1. Bernd Johann, Museum fur Nachrichtentecknik, Germany

3764A Opt.001 http://www.museum-nt.de/objekte/hp_3764a.html
3779C http://www.museum-nt.de/objekte/hp_3779c.html

2. Hugh Walker

3710A MLA Generator with 3716A Plug-in
3708A Noise and Interference Test Set
3709A Constellation Display
3746A Selective Level Measuring Set
3754A 25 MHz Access Switch
3762A/63A Pattern Generator and Error Detector (PO Tester 273/274)
3764A Digital Transmission Analyzer (Std)
3770B Telephone Line Analyzer (PO Tester 56)
3776A PCM Terminal Test Set
37722A Digital Telecom Analyzer
37711A Digital Data Tester
37724A SDH/PDH Test Set with 37772A Plug-in
37907A Signaling Advisor
70842B 3 Gb/s Error Detector Module
J2127A Transmission Test Set (Cobra)

Most of these units in private collections also have service manuals.

Appendix 5

Notes for Collectors and Restorers

This section contains information and a few tips for those interested in collecting and restoring instruments manufactured at South Queensferry.

Serial Numbers

Until the late 1990s, South Queensferry used the HP standard format for serial numbers, which gives a quite a lot of useful information on the age of the product.

Here is an example of the serial number on a 3776B PCM Terminal Test Set:

2423 U 00526

The first four digits denote the time of manufacture. The first two digits, 24 in this example, denote the year of manufacture. HP used 1960 as the reference (zero), so to get the year of manufacture, you need to add 60 to the first two digits, which in this example gives the year as 1984. The second two digits are the approximate week number in that year when that particular batch of instruments was made, so in this case it would indicate May/June.

Next, the letter "U" indicates the product was manufactured in the United Kingdom. Another letter commonly seen on instruments is "A" which indicates it was made in the USA. Incidentally, this also applies to computers and calculators. I have an HP32S pocket calculator special edition produced to celebrate HP's Golden Jubilee in 1989. It has serial number 2913A12184. Here is a typical serial plate from a South Queensferry instrument.

Going back to the example serial number above, the last five digits indicate the unit number in the full production life cycle of the product. The main production cycle starts with serial number 00100. Apparently, this was done to avoid customers thinking they had been sold an extremely early production unit with possible teething troubles! So to get the actual unit number in the production cycle, you need to subtract 100 from the serial number. In the example, that would be 426. If you get a unit with a serial number 0005X, it means that it is a production prototype. Usually six to ten prototypes were built, 51,52, 53 etc.

Some early units from the 1960s did have a slightly different format. A 3702A MLA Receiver has the number U825-00281. U means UK as before, and in this case the year number is 8 for 1968, with the week number 25.

Service Manuals and Servicing

For the collector, the service manual is as valuable as the instrument itself. At South Queensferry, a great deal of effort went into the preparation of the service manual, with beautifully drawn circuit schematics and board layouts, specification performance checks and descriptions of how the circuits worked. There was also a complete list of component parts for each assembly, all contained in a sizeable ring binder. Only with the service manual can you fully appreciate the engineering design that went into these products.

Up until the late 1980s, all instruments were shipped with a service manual and also an operating manual. Because of the size of these publications, service manuals are rarely available from on-line sources, though it is always worth checking. More of the operating manuals are available, including some on the Agilent website under “Obsolete Wireline Communication Test Equipment”, and “Document Library” at:

http://www.home.agilent.com/en/pc-1000003074%3Aepsg%3Apgr/older-obsolete-wireline-communications-test-equipment?nid=-536900202.0.00&cc=GB&lc=eng

The service information was intended for repair of instruments down to component level in a service office, and mostly the construction was conventional through-hole soldered printed circuit boards using industry-standard components. These are mostly still available to approximately the original specification, so the instruments are fully repairable by the knowledgeable collector. The only exceptions are where proprietary HP components were used such as ICs, hybrids and of course Read Only Memories containing firmware. If any of these are faulty, it is necessary to find a donor unit.

After around 1990, almost all new South Queensferry products used Surface Mount Technology (SMT) on the printed circuit boards. These were no longer intended for repair down to component level as SMT is too difficult to work with in a conventional service office. The approach then was board/assembly exchange when repair was needed. Sadly, this meant the instruments no longer went out with all the circuit diagrams and so are rather less interesting for the electronics enthusiast. The best one gets after 1990 are the specification performance tests, a block diagram and a process for fault finding down to assembly level. When preserving one of these instruments in working condition (mostly Digital Transmission Analyzers from the 1990s), it is best to have two units so that one can be used as a donor.

Finally, here are a couple of tips about products produced in the 1980s.

Keyboards

Most of the products from this period (before we went to membrane keyboards) used push button keys with a red plastic module having contacts onto the PCB. These modules were thermally welded onto the PCB with two spigots through the PCB. On the whole these still appear to work reliably, however over the years the action of the key has become stiff and puts up quite a lot of resistance when you push it. It doesn't feel good. By disassembling the front panel and accessing the keyboard, it is possible to see that each of the switch modules has a leaf spring that slides forward and then bows up as the key is pressed. The excess resistance seems to come from the spring failing to slide forward over the plastic housing. By applying a very small amount of silicone grease under this spring, the problem is greatly reduced. To do this, press the key down fully so the leaf spring is bowed up and then apply a small amount of grease under the spring with a fine wire bent like a hockey stick. Do this gently so as not to displace the spring as it is awkward to put it back with others keys round about. If you need to remove a key switch for cleaning or repair, just remove the thermal welds on the back of the board with a craft knife, and the switch module should drop out. It can be refitted using a small amount of epoxy adhesive (e.g. Araldite) in the spigot holes.

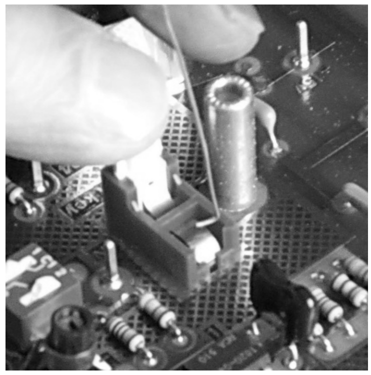

Battery Back-up

Some of these products use a rechargeable battery to keep alive some memory in the instrument so it retains settings and measurement results when the power goes off. This was considered a valuable feature for long-term tests, so the instrument would power up again and continue the test, having flagged the power loss.

The problem is these batteries are now getting old and have started to leak, particularly if never replaced. **The corrosive liquid can be very damaging to the circuit boards and components and may destroy an otherwise good instrument.** This picture shows the damage caused in a 3746A SLMS due to battery leakage. Another vulnerable product is the 3764A Digital Transmission Analyzer which has a similar battery mounted on the processor board. Any batteries like this should be removed as soon as possible, and the instrument operated without the back-up facility

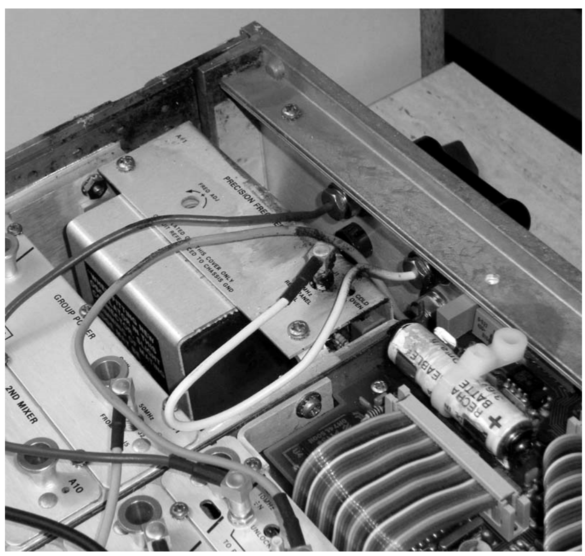

Appendix 6

Some Things that Might Have Been?

by Dr. Tom Crawford

Tom Crawford joined HP South Queensferry in 1966 and worked initially on the Microwave Link Analyzer. In 1967, he left HP to do a PhD at Heriot-Watt University in Non-linear systems measurement using multi-level Pseudo Random Sequences. Although not directly related, this research led to a better understanding of the properties of the Pseudo Random Binary Sequence (PRBS) which was central to much of Queensferry's test equipment for digital communications. After returning to HP in 1970, Tom worked on the Division's first high-speed Error Detector (3761A), and later went on to lead the R&D group which designed many of Queensferry's digital communications products in the 1970s and 80s. Tom and his team developed a number of innovative and highly successful products as described in Chapters 8, 9 and 10. These culminated with the "Modular-Hornet" program of the 1990's. The Company philosophy encouraged us to look for new measurement solutions – "contributions" which moved the industry forward rather than replicating products already on the market (which we called "Me too"). However, some of the measurement ideas were so "novel" that they never got beyond the prototype stage for various reasons. In this Appendix, Tom Crawford describes two "might have beens" developed in the 1980s, that never made it to the market. Tom (right) is seen here with John Doyle and John Young (CEO), early 1980s.

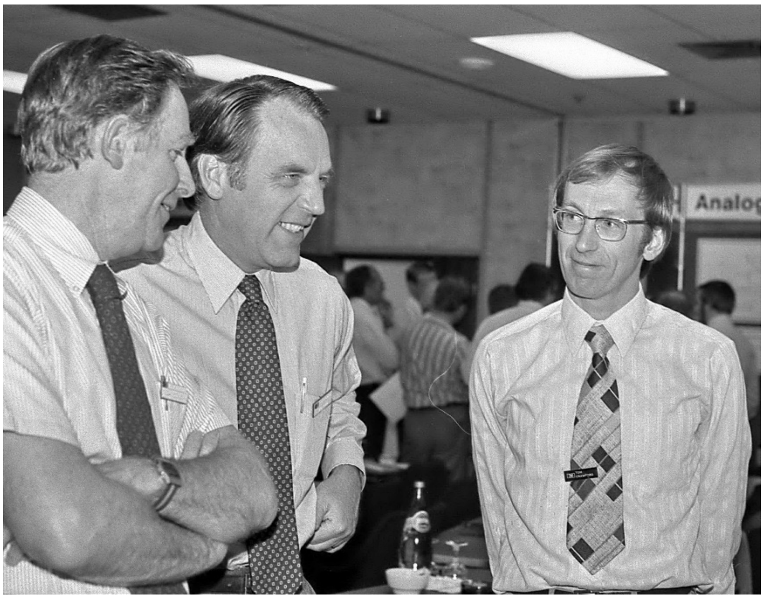

3786A Regenerator Margin Analyzer

Commercially successful R&D is always a balance. Too much "me-too" in the product line gives a blunt selling edge and a low return. Too much high risk, "into-the-unknown" R&D, leads to an empty pocket. Things which are too difficult can take so long that the market moves on to different needs and the window is missed.

Such was the fate of the 3786A Regenerator Margin Analyzer, with definition release in September 1980 and US Patent granted in 1983[1]. This product was targeted at metallic-cable PCM regenerators, usually at the primary rates of 1.5 Mb/s (T1) or 2 Mb/s (E1), carrying 24 or 30 telephone channels respectively, as described in Chapter 8.

The systems operated over twisted-pair copper wires that had each previously carried one analogue telephone channel between exchanges. Many copper pairs were combined in a single multi-pair telephone cable, with one pair used in each direction to carry the primary rate PCM signal.

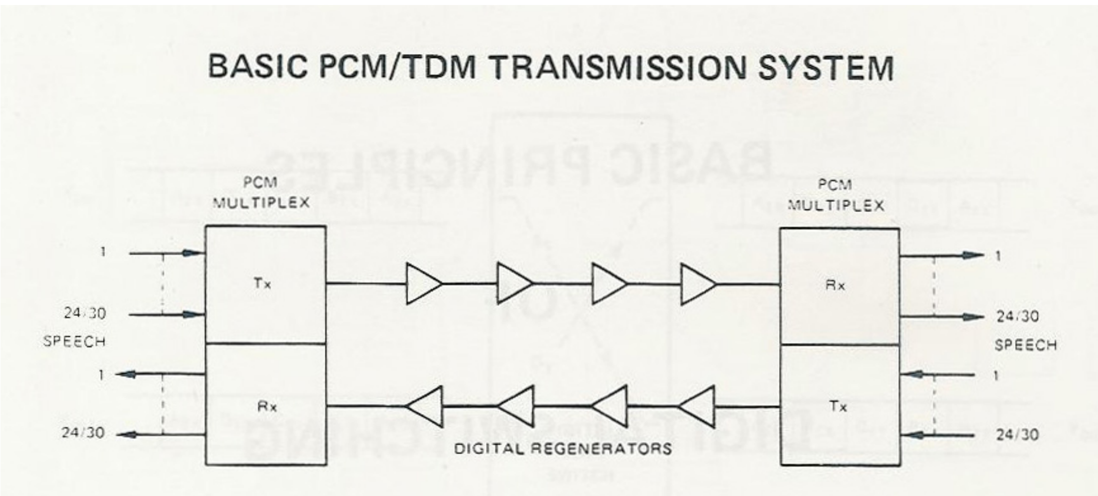

Typically, the 2 Mb/s regenerators were required every 1800m (1.2 miles), depending on the quality of the telephone circuits. The regenerators needed to compensate the significant high-frequency attenuation in the old telephone cables and also be able to operate error free in the presence of noise and crosstalk from adjacent circuits in the cable. In the early days of digital telephony, large numbers of these regenerators were required. In 1980, the British Post Office alone was ordering around 25,000 bidirectional units a year.

The classic approach to estimating the error-rate of a digital regenerator was to add thermal noise to a cable-attenuated version of a known signal and then measure the error rate produced. It was a realistic test since the crosstalk on the installed system resembled random noise, so the test checked that the regenerator had sufficient margin. The problem

[1] US Patent 4,384,354 "Noise Margin Measurement and Error Probability Prediction", Thomas Crawford, Alastair Reynolds and Ivan Young, HP Ltd. South Queensferry, granted May 17 1983.

was that a test relying on the statistics of random noise required a long measurement time which was a problem for high-volume production.

The novelty of the 3786A was to replace slow statistical measurements with fast deterministic measurements, so cutting final test times in regenerator production. The 3786A definition contained a fully featured product for designers and a simplified product for regular use in production-test. To achieve the deterministic measurement, a pseudo-random binary pattern sequence (PRBS) was encoded with the appropriate transmission line-code (such as HDB3 or 4B3T)[2] before being distorted by a cable simulator and fed to the regenerator input. The trick was that an individual pulse could be modified in amplitude when it was still an isolatable pulse before entry to the cable. This individual pulse could be modified in size until it introduced an error so revealing an effective threshold within the regenerator. By repeating this for all the various pulse positions, pulse polarities and trajectory directions, all the elements of inter-symbol interference could be taken into account.

This produced histograms of effective decision thresholds. These could be used to calculate error rates for a given probability distribution of additive noise. Problems arose when trying to match classic measurements to these new predictive methods. (This match was required because manufactures did not want to get into an acceptance fight with customers still using the existing method.) Second-order dynamic effects within the (unknowable) different designs of regenerators probably played a part in introducing discrepancies. The comparisons for HDB3 line-coding were reasonable but as the industry moved to new designs using more complex 4B3T, bigger discrepancies were revealed.

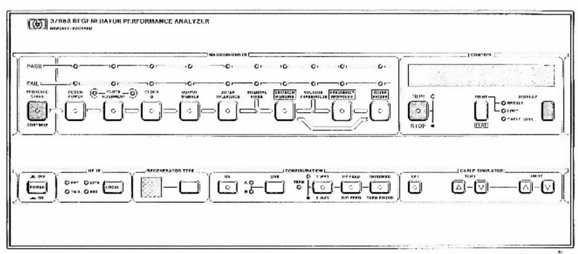

The project team with the prototype and a front panel drawing of the proposed instrument

The final straw was that by the mid-1980s interest in metallic regenerators was on the decline as the new fibre optic systems began deploying more rapidly. Continuing to work

[2] These are "ternary codes", meaning they have three voltage states, "+1", "0", and "-1". HDB3 is a simple pseudo-ternary code that inverts alternate marks (1s) to create a bipolar signal with zero DC. 4B3T is a true ternary code that converts four binary digits to three ternary symbols, thereby achieving a reduction in transmission bandwidth. The relationship between the binary data and the 4B3T line code is more complex.

on solutions to more and more subtle problems involving metallic regenerators was no longer an attractive way to spend R&D money.

However, there is a very interesting coda!

The HP South Queensferry US Patent which underpins the 3786A has been cited more or less every year from 1987 up to the present day. (Cited in 34 US patent filings between 1987 and 2012.) Companies citing our patent over this period include Lucent Technologies, IBM, Wavetek/Wandel & Goltermann, Siemens, Alcatel, Nippon Electric and Rambus Inc. who have been the most active since 2005. Some of these involve "drift tracking feedback" inside communication systems. From the test equipment side, the W&G citing is interesting with a "Non-invasive digital cable test system" in 2002. Maybe the 3786A should have leapt into an optical form[3] after all?

ALBERT

The ALBERT investigation started under the title, "An IC for Low-speed Bit Error Rate Testers" with a memo issued by Virgil Marton and myself in January 1985. Whilst the 3786A was to be a new solution to an old problem, ALBERT was to be a new solution to an emerging problem. The problem was the growing diversity of formats emerging in overhead and payload structures in digital communications applications. Traditionally this meant continuously reinventing hardware and the software to control it. ALBERT was conceived as more general solution, structured in three distinct sections for ease of definition and modification. By 1987, the architecture was defined and a team of six engineers were working on the project.

A small number of configurable hardware "devices" would be designed once-only for each type of conveniently separable operation required to build the serial stream. These devices would be orchestrated by a specially designed "formats" processor to output their components into the signal stream. The top level software would drive the formats processor from a high-level graphical description of the signal.

In January 1990, five years after the project began, the team (by then 11 engineers) issued a detailed description of the project and the various technical elements, stating the objective as:

"The main objective of the ALBERT project is to develop an architecture which combines hardware and software elements to provide a programmable means of structured serial data generation, reception and measurement. The project will develop and validate the

[3] South Queensferry had a continuing interest in deterministic methods of predicting low error rates. The very low background error rates in high-speed optical systems were impractical to measure using conventional tests, and some application software was developed to work with the 71612 12 Gb/s Error Performance Analyzer to predict residual error rates from digital eye opening (see Chapter 10).

architecture, integrate hardware elements and provide a supporting software development environment to facilitate its use in the next generation of telecommunications test products."

Virgil Marton and David Grieve initiated a draft patent application as co-inventors of the ALBERT architecture. As David Grieve recalled, "The patent application was never pursued when it became apparent the chipset would cost $250k (we got as far as a mask set done by the Wolfson Institute). It didn't help that the Division was short of money at the time."

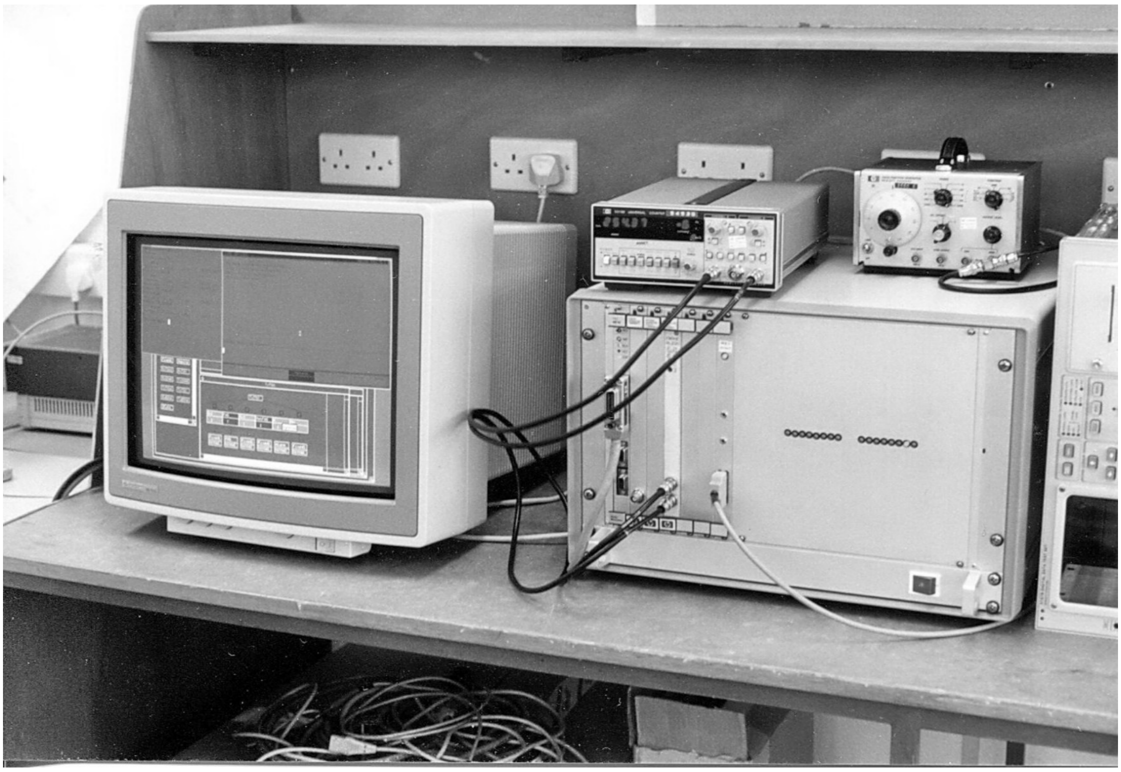

Again, the difficulties inherent in genuine innovation took their toll by extending the time and cost of development. During the five-year development period, improvements in commercial general purpose programmable logic (the Field Programmable Gate Array or FPGA) continued. The bit-rate limitations of ALBERT in the face of the upward trend in the needs of the market, blunted the appeal of ALBERT and it was discontinued in early 1990.

Alas there is no known interesting coda to tell at this time, although some of the knowledge and experience gained on ALBERT probably did benefit the new family of Hornet products of the 1990s which made extensive use of programmable logic, as described in Chapter 10.

Appendix 7

Instrument Serviceability and the Challenge of the 3745 SLMS

by Garry Irvine

Garry Irvine joined HP South Queensferry in 1970 as a product support engineer, having previously worked with Ericsson and Siemens. After initially working on the Microwave Link Analyzer, he joined the SLMS support engineering team and here he describes the special challenges of that new instrument. In 1980 he moved to Colorado Springs joining the team supporting Queensferry's products in North America. Later, he contributed to the new SS7 signalling test area and in the 1990s got involved with SS7 sales and systems development in the USA and Canada, culminating in work with HP's computer business and its interests in the telecom industry.

The tradition, and a recognised and appreciated feature of HP instrumentation, was the detailed user and service documentation. Traditionally HP products included in-depth descriptions of the technology being tested, as well as the various circuitry and functionality in order to simplify component level repairs. This was designed into every HP product to help reduce the "total cost of ownership" and provide "life-time support".

However, this tradition was at risk when the first of a new breed of processor-controlled instruments came along in 1975 – namely the HP 3745A/B SLMS. It was recognised early on that the repair of this product was going to be a huge challenge for HP's Field Service Organisation and customers who expected to repair HP's instrumentation down to component level. HP's field service engineers became familiar with the more traditional instruments because of the large volumes they saw of frequency counters, multi-meters, oscilloscope etc. In our case, Telecom Test equipment was almost always "new" as our shipments were relatively small and the reliability relatively high.

As with all of South Queensferry Division's Scottish products a support engineer was assigned to the R&D team. His or her function was to address the serviceability, manufacturability and maintainability of the finished product. In this case two engineers were given the daunting task of making HP's first processor controlled instrument serviceable! To meet this challenge the HP 15585A Diagnostic Kit was developed.

The troubleshooting strategy adopted was familiar to most engineers – that of simulating control to aid the finger pointing exercise during troubleshooting. The service strategy also included adding test links at strategic circuit positions to allow identification of faulty components for replacement. The bigger challenge was to educate the field engineers and customers, who had never seen and often never heard of central processing units or CPUs. Yes, the brave new world of computing - as we know it today - was starting. This meant that the test kit itself had to be described in detail along with Binary Coded Decimal (BCD), Octal number theory and machine coding for the processor being used which was the 8-bit Intel 8008 chip.

However, to help avoid confusion with something so new, the troubleshooting approach was always one of "cause and effect" - that is, switch setting "XYZ" would give results "ABC". The kit included a number of test boards that would simulate and emulate the various control functions of the micro-processor. The kit also included a Processor Test Assembly that was used to exercise the actual processor in order to determine if the processor board itself was malfunctioning.

Using the Diagnostic Kit was pretty involved because the instrument's keyboard assemble also housed the processor circuitry. So, using the Processor Test Board meant taking the instrument apart to allow board level diagnostics to take place. With such an unfamiliar instrument as the HP 3745A we built in a test mode that allowed the complete keyboard and processor to be set to "normal" mode in order to check that any repair done was effective. This was appreciated by Field Service Engineers, as it saved time in the assembly and disassembly of such a large complex instrument.

As far as the HP 3745 SLMS was concerned, meeting the serviceability goals turned-out to be invaluable to Hugh and friends who managed to bring one of these instruments back to life for preservation almost 40 years later!

Appendix 8

Developing a Test System for the CANTAT 2 Trans-Atlantic Submarine Cable

by Bill Lauchlan

Bill Lauchlan joined HP South Queensferry as a test engineer in February 1967. Like a number of the early employees, he came to HP from Ferranti in Edinburgh where he had worked on the testing of radar systems for the new generation of fighter jets in the 1950s and 60s. In the early 1970s he worked on the MLA production line, and moved to Loveland, Colorado a few years later to provide support for Queensferry's products in North America. Later he joined the sales organization in Canada for the rest of his HP career. Bill describes an interesting system developed in Canada and based round the 3746A Selective Level Measuring Set (described in Chapter 5), which automatically monitored the supervisory tones from approximately 500 submerged repeaters[1] on CANTAT 2 submarine cable. CANTAT 2 was commissioned in 1974 and operated until 1992, carrying 1840 bi-directional telephone channels with 3 kHz spacing, on a dual-band 12 MHz coaxial cable. The cable ran from Beaver Harbour in Nova Scotia to Widemouth Bay in Cornwall, a distance of around 3000 nautical miles (3500 miles or 5500 km). The system was a joint project in the early 1970s between the British Post Office (BPO) and the Canadian Overseas Telecommunications Corporation[2] (COTC, later Teleglobe Canada). The cable was laid in 1973 by CS "Mercury"[3]. Here is Bill's story.

[1] The transistorised repeaters (wideband analogue amplifiers, probably designed and built by STC Greenwich) were spaced every 7 miles (6 nautical miles) and were housed in a cylinder 2.5 m long and 250 mm in diameter. They could withstand pressures of more than 4 tonnes per square inch on the ocean floor. The 500 repeaters were all powered through the coaxial cable from the shore-ends using a high-voltage supply, typically 10 to 15 kV.

[2] COTC handled all international telecommunications across the Atlantic and Pacific by undersea cable or satellite.

[3] Cable Ship "Mercury" was built for Cable & Wireless by Cammell Laird in 1962 and was finally scrapped in 1997, having laid a number of the major trans-oceanic cables.

This was a really interesting project on many fronts.

In around 1985-1986 COTC was faced with obsoletion of the Test/Troubleshooting System for CANTAT 2. This old equipment, from Marconi or BPO, was called the SRME (Submerged Repeater Maintenance Equipment), but was no longer supportable.

On a routine visit with Teleglobe in Montreal along with Gerard Joly (Field Engineer), we met with a Submarine Cable Engineer (Harry Bleeker) who had responsibility for support of various Undersea Cables from Canada to Europe, the Caribbean and under the Pacific to Australia and New Zealand.

Harry presented me with the problem, asking if we could come up with something to do the job.

The CANTAT 2 cable had around 500 repeaters under the Atlantic between Canada and the United Kingdom. The multiplexing was FDM. I forget the FDM Plan but it was High-Band one way and Low-Band the other way. There were some "empty" spaces in the spectrum which were used for Supervision and Testing. At the voice level it used "TASI" (Time Assigned Speech Interpolation), a technique of re-using channels during the statistical 50% of time when one party in a conversation is not talking. This probably increased capacity to around 4000 simultaneous calls.

The SRME equipment carried out a number of functions:

1. Finding breaks in the cable.
2. Identifying the 500 repeaters
3. Noise measurements
4. Tone measurements.

Most of these tests were Loopback Measurements, whereby a signal would be sent from Beaver Harbour and, depending on its frequency, would loop-back at a particular repeater and then be transmitted back to Canada where it would be measured.

To make this happen, each repeater had a Loopback and Frequency Shift Circuit in it, consisting of bandpass filters, a mixer and a crystal local oscillator. The crystal oscillator for each repeater had a different Frequency.

After appropriate investigation, I said we would come up with a system design to address the application.

I had in mind to use the HP 3746A SLMS, HP 3336A Level Generator and a Controller, as the main hardware components and develop software using some standard tools like HP ITG (Interactive Test Generator) and structured programming methods as per HP TAIPAN architecture. ITG was a software development tool created by HP in the 1980s to help Test & Measurement people who were "non-computer types" write effective control programs more easily using libraries of sub-programs and drivers[4]. The Controller we ultimately settled on was the HP 9836A with nice big screen and graphics.

[4] For further information on HP ITG see http://www.hpmuseum.net/display_item.php?sw=438

It's hard to describe this project without getting into lots of internal HP politics involving the integration of HP Technical Computer Group and Test & Measurement, an ill-fated venture which became known as "The Dark Ages" to many of us who were around and survived that period! Suffice to say, it caused me some major nightmares getting this job done. In due course, we got a contract to develop and supply this System.

I was the design authority and I suppose Project Director, Manager and general "Dog's Body" who did the things nobody else had assigned to them. I had some folks from our Computer Group Project Centre in Toronto to help carry out the software development. Much of my time was spent trying to get this project moving ahead, and avoiding "analysis paralysis" in our Project Centre.

However, we did manage to get the job done.

Once the System was developed, a deputation from Teleglobe visited Toronto to go through acceptance testing at our facility. The System passed with "flying colours" in our lab, using simulated signals since we didn't have a real Undersea Cable there. The customer was real happy.

After initial acceptance, the equipment was all shipped to Beaver Harbour for installation. I went down to Nova Scotia to carry out installation and turn the system up. This happened just before Christmas, and that was when the fun began!

The equipment all installed just fine and we proceeded to testing it on the real live undersea cable. The various measurements for noise and gain all worked fine, but one routine for identifying each repeater and locating breaks was totally useless. This was probably the most important measurement of all. The reason it was useless was 'cos I decided to use the Counter capability in the 3746A SLMS, whereby once the Repeater ID tone had been detected the counter would be turned on and it would give the exact frequency. This test had been run in our lab environment with simulated signals and worked really well. The difference with the real cable was poor signal-to-noise ratio (SNR). The counter couldn't make its mind up!

When Harry Bleeker saw this, he looked like he was going to have a heart attack and that his whole project was going to fall down around his head. I managed to calm him down and, as is usually a good idea in these situations, I took him for supper and a pint. Over supper and an ample supply of napkins, I explained to him what the problem was and that we would fix it. At that stage I didn't know how, but I knew we would. He didn't seem to have the same level of confidence.

The way the Repeater ID worked was that each repeater had a unique local oscillator frequency. These nominal frequencies were all held in a database in the HP 9836A Controller.

However, due to variations in ocean temperature, these frequencies would drift by as much as +/- 50 Hz (maybe a bit more). This drift was somewhat predictable. So, using the narrow flat-topped 22 Hz filter, it was necessary to scan a little bit round about each nominal frequency to find the tone and then get a steady reading. With the poor SNR, this took forever. I think we actually completed one scan of the whole cable and it took about

12 hours. This was all supposed to happen in about 20 minutes. Something radical had to be done to make this work.

So, we got busy with the thinking caps and the napkins.

I recalled a terrestrial cable measurement setup from my time in Loveland, Colorado, where AT&T used the HP 3745B SLMS for what I think were called PEARL Tones (Performance Evaluation at Remote (Repeater) Locations). In the AT&T setup, I don't think there was loopback, just a tone generated at the repeater. I think this made use of the "Picket Fence" effect, if you choose the right step-size when scanning the spectrum. This can allow you to determine the frequency quite accurately.

By this time it was about a week before Christmas and I was heading off to Scotland direct from Nova Scotia. We wrapped up our work at Beaver Harbour, and Harry headed back to Montreal for Christmas. I told him we'd have a fix when I got back in January and I set off to Scotland.

On arrival, I paid a visit to South Queensferry and prevailed on the Service Department (Bill Morrison I think) to let me borrow a 3336A, 3746A and a 200 Series Controller over the holiday period which I took home with me to Comrie.

The new technique I experimented with was to take three readings with the correct step size so that I would get at least one full-level reading on the flat top of the filter. This meant there would also be at least one on the upper or lower skirt of the 22 Hz filter. By having the filter characteristic stored in smallest increments possible (I think it was 1 Hz) and using this as a look-up table to compare the measured readings, I would be able to determine tone frequency with great accuracy. That was the theory and it worked like a charm. By only having to take three readings it was also very fast, and also no counter was involved.

Earlier I mentioned using HP ITG and the TAIPAN architecture from Corporate Systems Engineering Organisation (SEO) folks. This was an effort to get some order into software being developed piecemeal throughout SEO on the 200/300 Series Controllers and also to have a library of software routines available for re-use. I don't think some of the major objectives were accomplished but, at least from my perspective, TAIPAN did produce a good Structured Programming methodology for the Systems Engineers, which I certainly took advantage of. I now recall the name of the main TAIPAN promoter in SEO. It was Jon Kim, I believe.

The reason I diverge to mention this, is that through having used this architecture, it was very easy for me to just throw away one module in the software we had developed, and replace it with another using the measurement technique described above. This kinda blew Harry's mind away when I got back to Nova Scotia and basically plugged in a new software module to solve problem.

On installing the new software in Nova Scotia, we tested the new method of identification and it worked brilliantly and real fast. I know it beat the 20 minutes, but I can't remember by how much. Unfortunately, we were not able to test the cable with a real break in mid-ocean at that time, but I'm sure the fishermen on one of the continental shelves, facilitated that at a later date! It is worth noting that there were NO after sales support problems with this System right up until the time the CANTAT 2 cable was decommissioned, to be superseded by digital.

In the context of our Project Centre fiasco, the last sentence is important because our computer folks would not adhere to a fundamental rule which I have always applied. That is not to take on a project unless you know both the technology and application very well. They always looked on the customer's measurement and/or data acquisition as just some kind of head-end feeding them data which caused them great grief and huge expense by having to sort out a big mess after their systems were installed.

As a footnote to this Project, I went on a holiday in 1987 using many Frequent Flyer Points to travel first class across the Pacific. Courtesy of Harry Bleeker, I was able to visit the Teleglobe Cable Terminal at Keawaula, Hawaii along with Jack Weldon (Sales Engineer), where we were treated like VIPs by the local Teleglobe Manager and I got to see what an undersea repeater really looked like up close[5]. I would have loved that Station Manager's job.

I always remember the direction he gave us to get there from Honolulu:

"*Just drive up the coast past Pearl Harbor until you run out of road, then turn right.*"

We also visited with the Telco in Fiji on that trip. Shortly after we left Fiji, the government and many officials, some of whom we had met at a Rotary meeting, were locked up in jail. A coup took place and I believe Fiji became a republic. It was re-admitted to the Commonwealth in 1997. I'm not sure if our visit to Suva had anything to do with that.

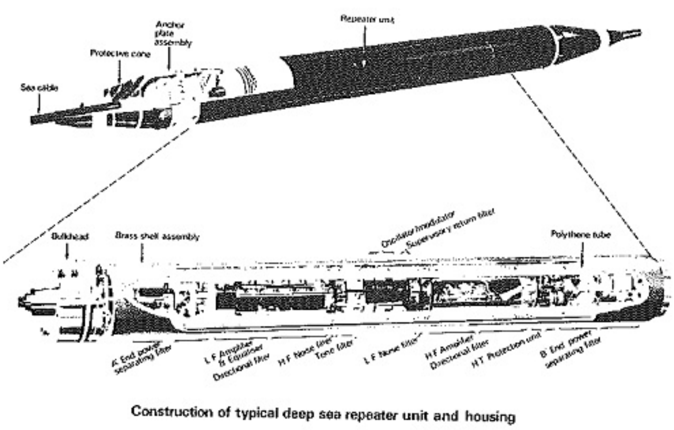

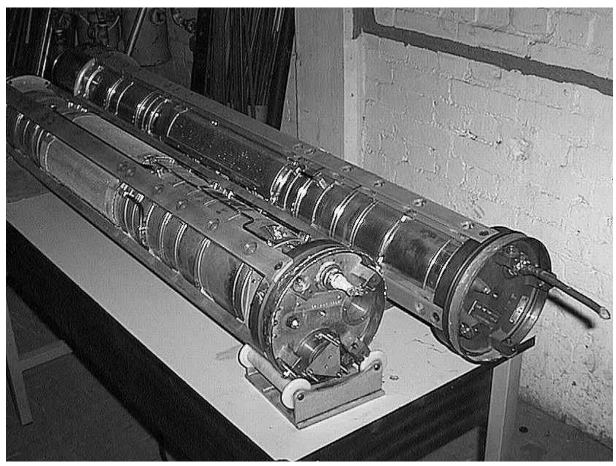

CANTAT-2 Repeater and a view of the insides manufactured by STC Greenwich

[5] For more on the design of repeaters, see https://atlantic-cable.com/Article/SA/65/index.htm

Product Index

General Index

UV

W

XY

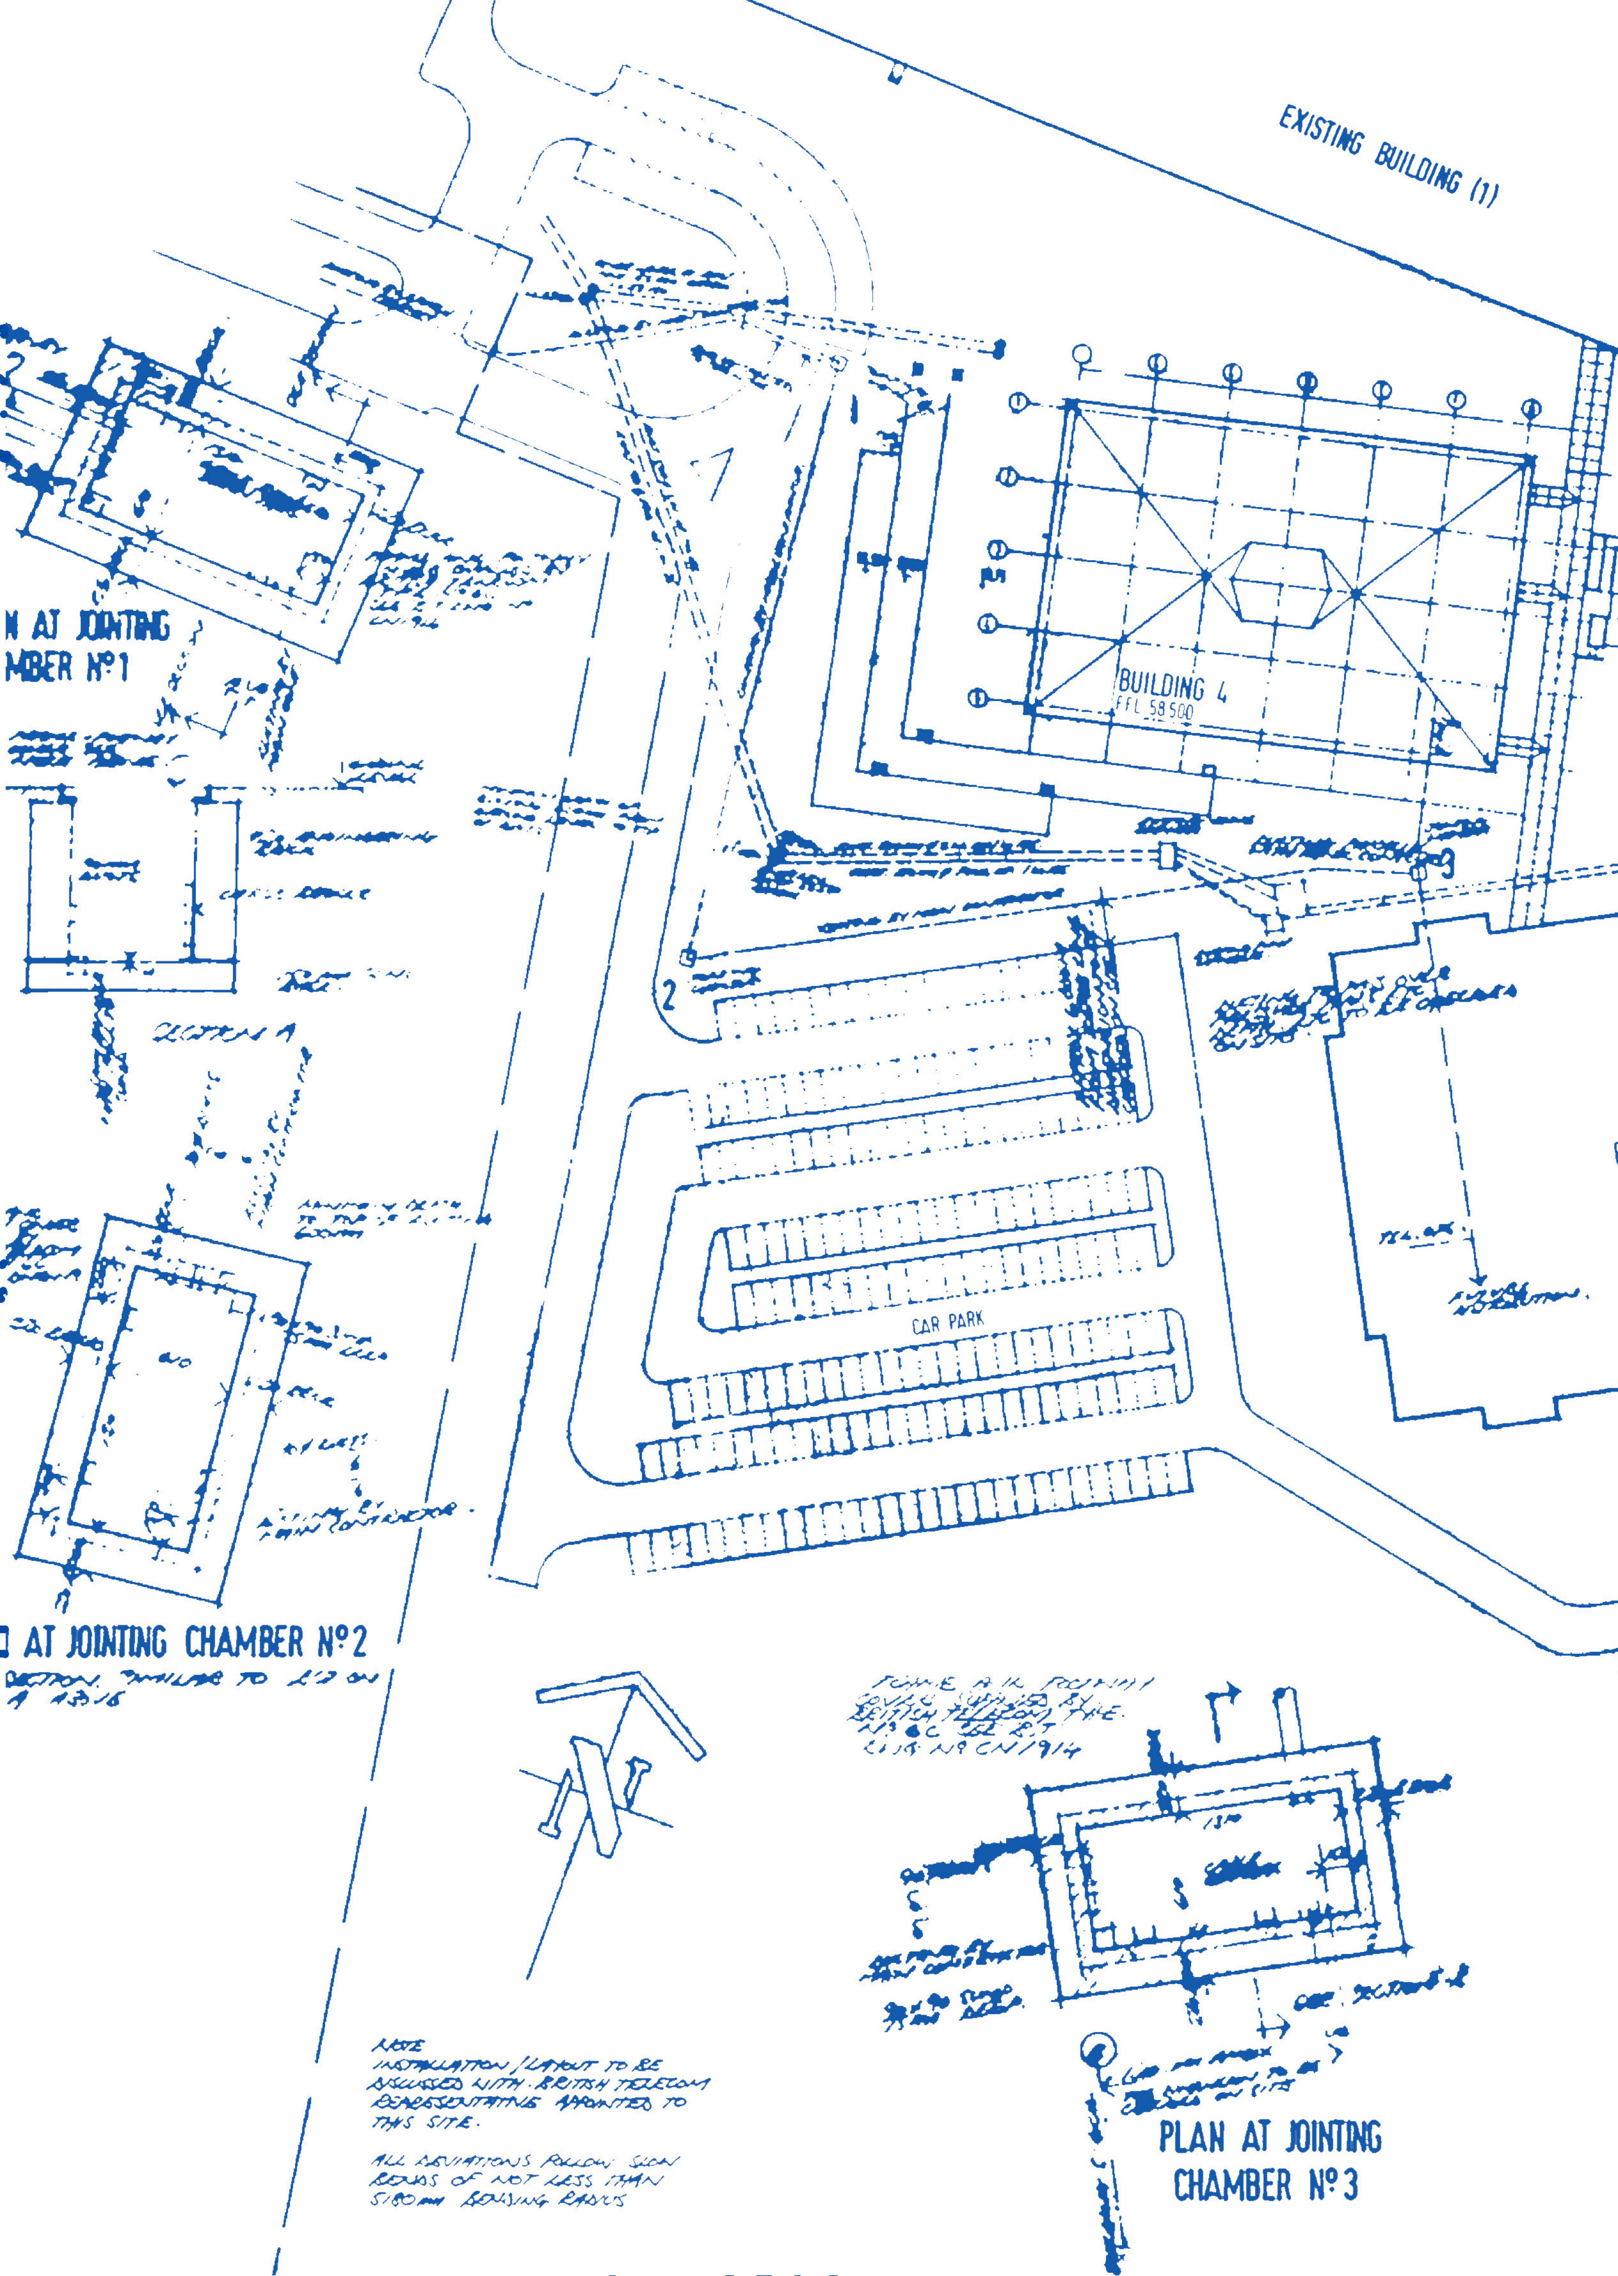
EXISTING BUILDING (1)
BUILDING 4
FFL 58500
CAR PARK
AT JOINTING CHAMBER Nº2
PLAN AT JOINTING CHAMBER Nº3
NOTE
INSTALLATION / LAYOUT TO BE DISCUSSED WITH BRITISH TELECOM REPRESENTATIVE APPOINTED TO THIS SITE.

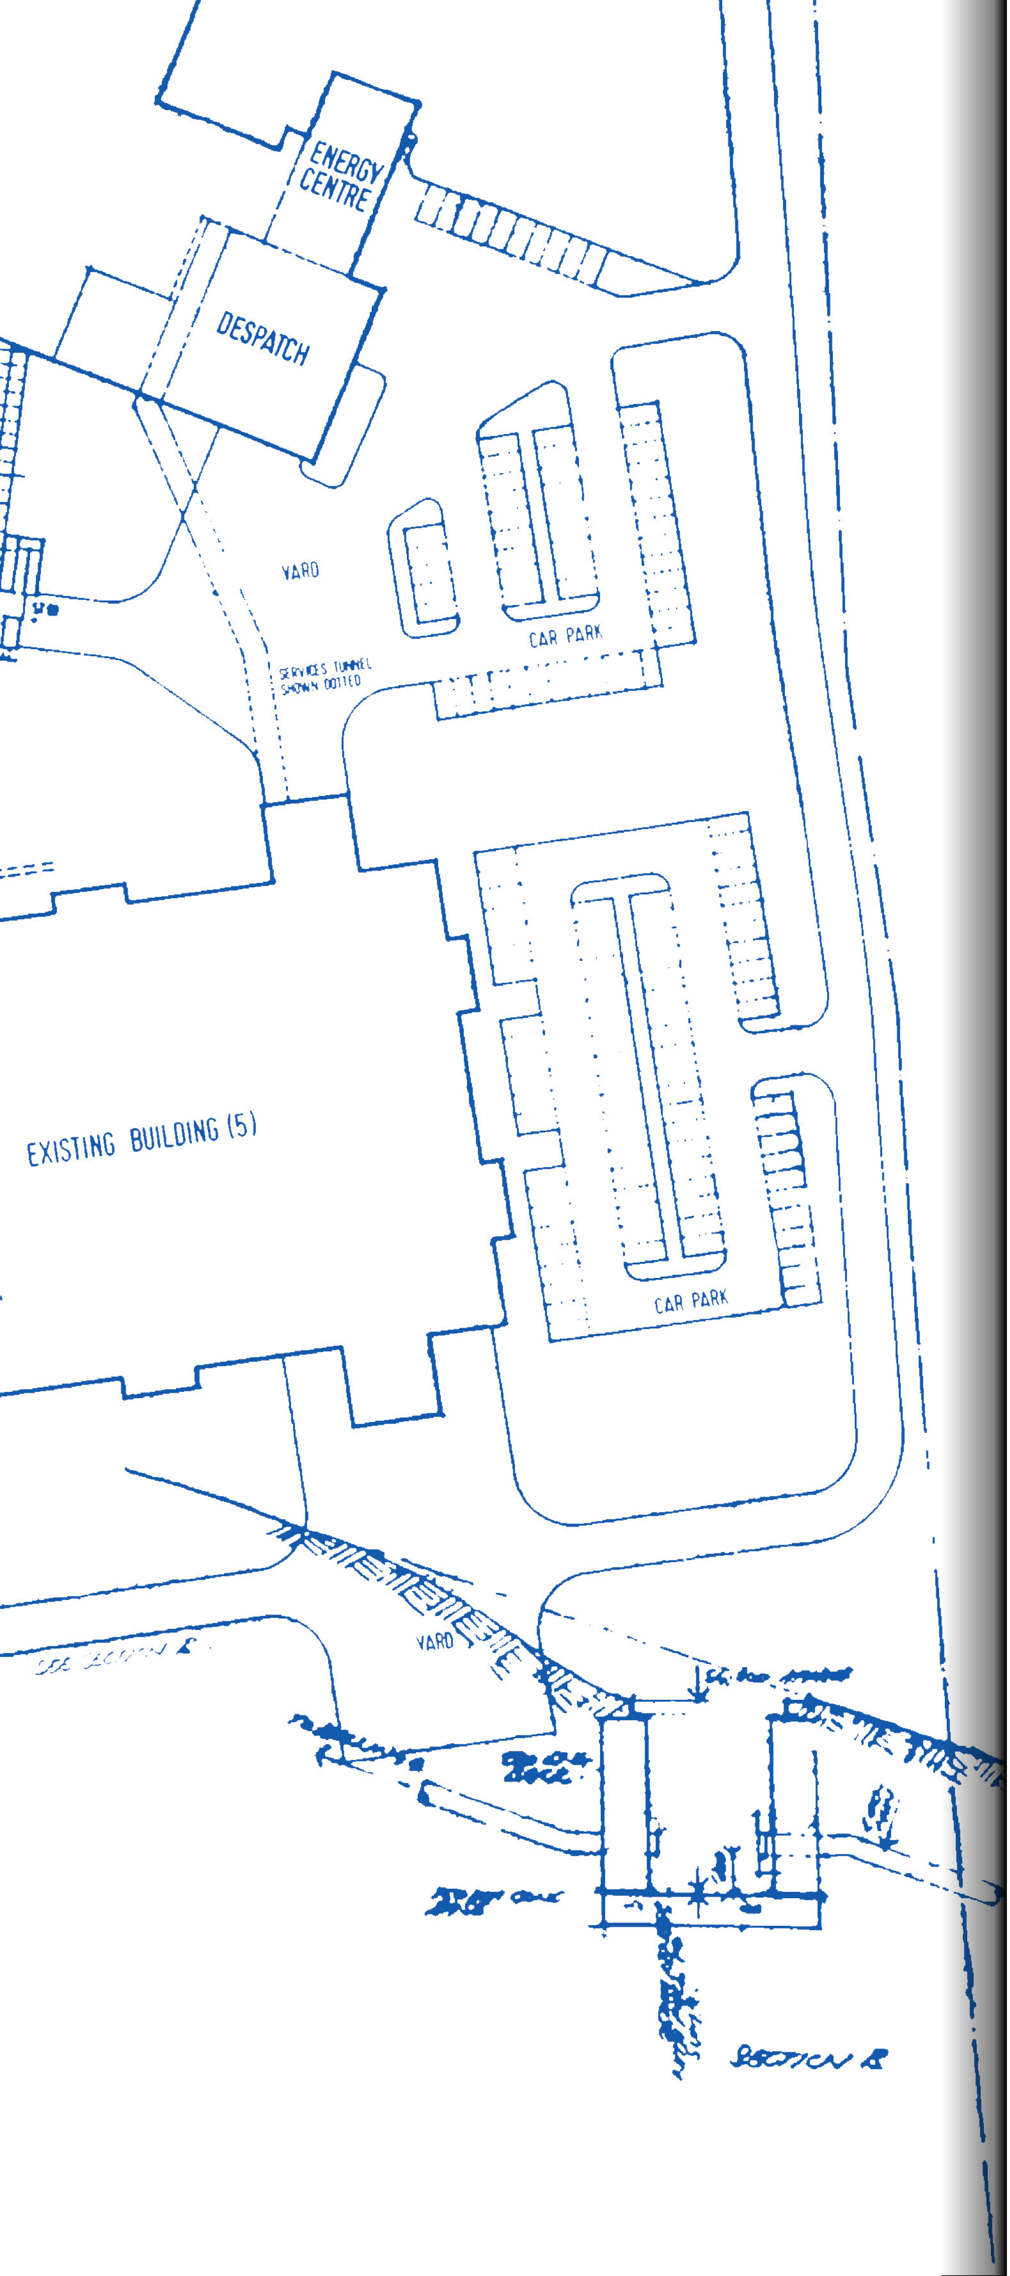

The Author

Hugh Walker was brought up in Coventry, before graduating with engineering degrees from Cambridge and Southampton Universities.

He joined Hewlett-Packard at South Queensferry in 1970, where over the next 31 years he held various design, marketing and management positions working in telecommunications test.

Hugh was in a unique position not only to realise that the story of "the factory" was worth telling, but in fact needed to be told.

He also had the contacts and years of service to give a fascinating insight into the rise and fall of one of the premier players in Silicon Glen.


HEWLETT PACKARD
37204 HP-IB EXTENDER
POSITION
PORT A FAULT
PORT B FAULT
POWER ON
HUGH WALKER
MARCOM MANAGER
Queensferry Telecommunications Division
hp
HEWLETT
PACKARD
South Queensferry
West Lothian EH30 9TG Scotland
Telephone 031-331 1000
Telex 72682
Telefax 031-331 3000
1987

03700-90002