Packet Switching: The first steps on the road to the information society

Abstract

Today's information-based society is built on the precept that information, sound and moving pictures can all be expressed in the form of data, transferred using electrical, optical or radio-based telecommunications. The concept of packet switching, which made this possible, is not well covered in the historiography, getting lost somewhere between the descriptions of advances in voice technology and the history of computing. We first examine the genesis of the idea, before focussing on an initial implementation in the form of the British Post Office (Telecommunications) (BPO-T) Experimental Packet Switched Service (EPSS). This examination exposes how rapid progress towards a useable standard could be made despite the primitive nature of the technology available. We will look at how the technology functioned and was used, before demonstrating how, as its popularity within business circles increased, it was overtaken by the Internet's approach to packet switching. We will show that as consumers became more computer literate, events conspired to shift from restricted access services based on infrastructure from a single accountable provider, to a more widely available set of facilities in an open networked environment, owned by many providers. Growth is discussed in terms of accepted historically based ideas from Kuhn and Winston. We make our case by bringing together the experience of practitioners involved in the delivery of these early packet based services and the available literature. Our contention is that without packet switching, today's information-rich communications services, which are vital to the global economy, would not exist.

Introduction

Packet switching underpins the Internet and many private networks and is a major contributor to the world economy, supporting many of the business and entertainment services that are fundamental to today's society. This paper identifies how the basic ideas were generated, by examining the development of the UK's first packet-switching service, which was an early step on the evolutionary path leading to today's information-rich connected world.

The early concepts that gave birth to this capability will be explained and placed in a historical context. By focussing on a particular implementation, the British Post Office (Telecommunications)'s (BPO-T) Experimental Packet Switched Service (EPSS) service, the importance of ideas and people in the development of this fundamental capability is illustrated. We then show how this service evolved to a service based on international standards, leading to the current connected world.

The paper will identify the first steps in developing packet-switched solutions, describing the technique itself, the origin of the ideas and the external factors driving it. Having established the engineering context, we describe the path from prototypes to EPSS and examine the system and its components in more detail. We then demonstrate the significant impact this system had on data communications, its relationship to parallel developments and its contribution to developing the modern Internet and the Information Society.

Why divide data into packets

In the 1960s, there were few specialist data networks. Communication between computers was usually carried across the public switched telephone network (PSTN) or through dedicated private circuits. The private circuit solution was cheaper for higher data volumes, but typically required point-to-point connections on a one-to-one basis. This was expensive and networks could be complex to manage.

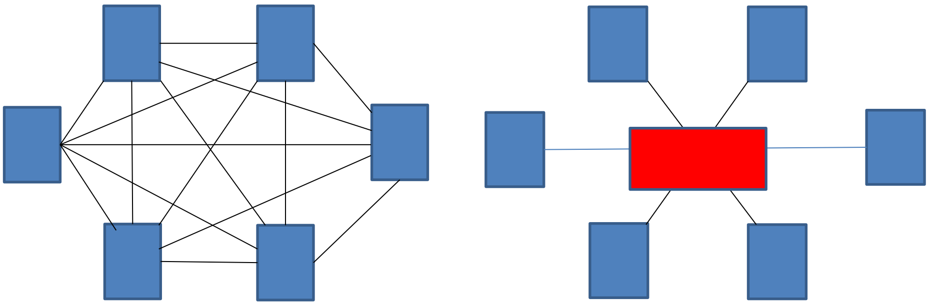

a) Connecting 6 computers (blue) in a mesh needs 15 private lines

b) Connecting the computers via an intelligent switch reduces this to six, saving costs

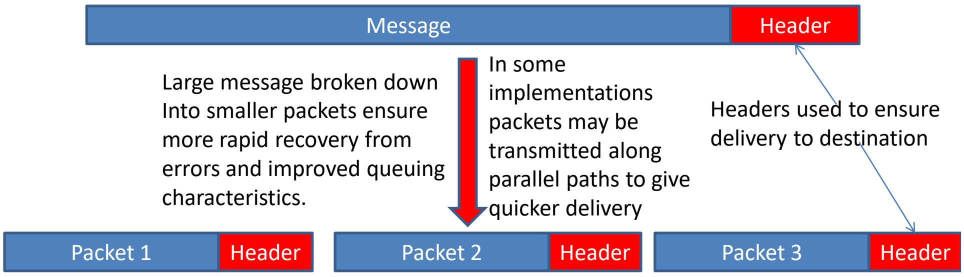

Large message broken down Into smaller packets ensure more rapid recovery from errors and improved queuing characteristics.

In some implementations packets may be transmitted along parallel paths to give quicker delivery

Headers used to ensure delivery to destination

c) Turning message switching into packet switching

Figure 1 – Evolving from a mesh to a packet switched network

One solution to this was packet switching, as illustrated in figure 1, which was independently devised by Paul Baran in the USA and Donald Davies in the UK [Abbate, 1999; Kirstein, 2009; Rutter 2005, pp 2-3; Naughton 2000, pp 92-95].

Baran, who worked for American defence contractor RAND, adapted store and forward messaging, a technique used in telegraph and telex systems. Message switching devices (or nodes) would temporarily store a message before passing it onto the next node in the chain connecting users, eliminating losses in signal strength with transmission distance. Message switches also allowed a single telex machine to connect to other telex machines on the network, based on the necessary route being set up between the two end points, analogous to setting up a call in the telephone network [Abbate, 1999].

Baran's innovation was to segment the message into smaller sections and route these based on information, known as headers, added to each packet. The message switches' role would be taken by a network of small and relatively inexpensive minicomputers, connected to each other using standard telecommunications links. Small packets would improve the queuing characteristics of the network, making recovery from transmission errors more efficient and permitting faster transmission though use of multiple routes [Kirstein, 2009; Roberts, 1978].

These ideas were driven by the needs of a defence-based organisation during the cold war and, since survivability was vital, the nodes could be sited away from centres of population that would be military targets. The routeing capabilities of the nodes would also provide alternative paths between network end points, reducing the impact of anticipated heavy damage [Abbate, 1999].

The UK government was concerned about the widening technology gap with the USA. Their focus was on exploiting the "the white heat of the scientific revolution" supported by the Ministry of Technology. Labour wanted to redirect research and development activities away from military projects towards civilian industry [Abbate, 1999], from civil aerospace and nuclear power to other sectors and to migrate from government laboratories to the private sector. The change was interpreted in terms of Britain, rather than not being innovative, lacking a strong industrial organisation, good management and attention to application [Edgerton, 2002].

Davies, a computer scientist at the UK's National Physical Laboratory (NPL), was interested in interactive time -sharing systems, which were limited by inadequate data communications. He knew of message switching in the telegraph system and proposed dividing messages into standard sized packets and having a network of small computers to route packets based on information carried in packet headers. This technique, which he called packet switching [Pelkey, 1988; Roberts, 1986] was expected to bring down costs. Following public presentation of his ideas, in late 1965 and early 1966, Davies became aware of Baran's work [Abbate, 1999; Pelkey, 1988].

Since an important stimulus to the EPSS work came from the work at NPL, we begin the assessment of the journey from concept to an implementable network, by starting with the NPL development.

The journey from the laboratory

In December 1965, Davies proposed that the BPO-T should build a prototype network; but NPL, which only had resources to build a small prototype network called the Mark 1; this is shown in figure 2. This prototype network was built between 1966 and 1969, providing internal services using 768 K bps channels and was operational by 1970 [Kirstein, 2009]. It covered the NPL site, using a single Honeywell DDP-516 minicomputer, as a packet switch [Kirstein, 2009]. Its development was described at a 1968 conference, two years before similar progress on ARPANET, the precursor to the Internet, was demonstrated [Kirstein, 2009, Roberts, 1978].

Davies would learn of parallel work going on under the US Department of Defense (DoD), funded Advanced Research Project Agency (ARPA) banner at a UK symposium and in 1967. By the end of the 1960s, ARPA had funded a variety of time-sharing computers located at universities and other research sites across the US, which ARPANET was to connect.

Roberts, who managed the ARPANET project, learned of Mark 1 at a 1967 symposium in Gatlinburg in the USA. He adopted Davies' term packet switching and some aspects of the NPL design, when considering the development of ARPANET. The consultancy company Bolt, Beranek and Newman Inc (BBN) was engaged to develop and build the network and Baran was recruited into ARPA in 1967. When it was implemented in December 1969 the ARPANET had 4 nodes [Roberts, 1978].

The UK government favoured collaboration with Europe rather than the USA [Kirstein, 2009, Kelly and Lee, 1978], so engagement with ARPANET was limited to Peter Kirstein's group at UCL. Kirstein highlights that the UK had a more unified approach to research funding, allowing a more integrated, but insular, approach to computer network development, lacking the impetus to the research that ARPA provided.

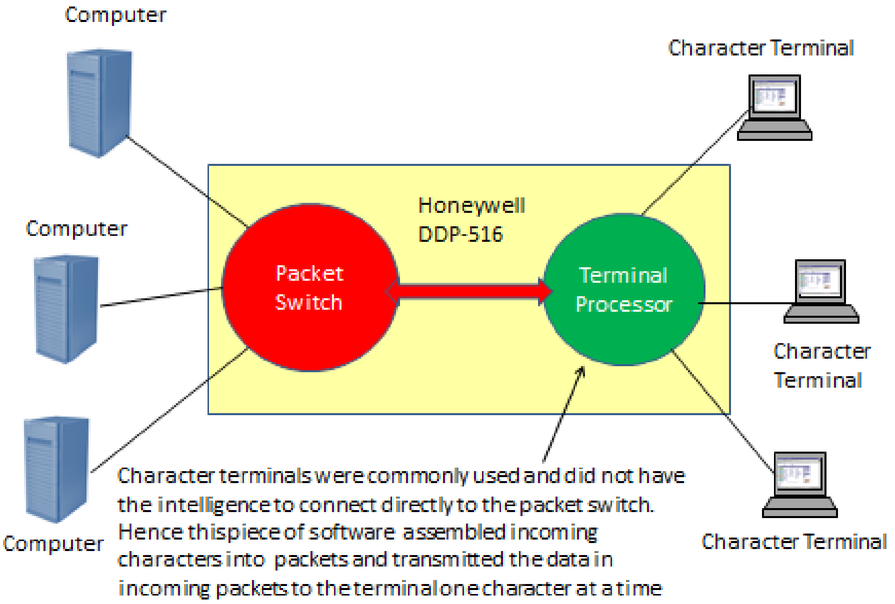

Figure 2 – Schematic of the NPL Mark 1

BPO-T staff attended a 1966 lecture by Davies, resulting in the BPO-T showing polite interest, but no serious consideration of packet switching [Pelkey, 1988]. The British equipment manufacturers were less interested; Baran and Roberts faced similar issues in the USA [Roberts, 1986].

Davies maintained good relationships with the BPO-T, who took an independent approach but were lukewarm at senior levels. Stanley Gill of Imperial College was a founder member of the Real Time Club; an organisation that exerted pressure on the BPO-T to improve its provision for data communications [Campbell-Kelly, 1987]. In November 1969, The Real Time Club submitted a proposal to the Postmaster General for an experimental network [Roberts, 1978, Campbell-Kelly, 1987].

This was an unsettling time for the BPO, which was managing the transition from being a Civil Service department to a nationalised industry and faced the spectre of the Carter committee and what it might recommend. Data, although an increasing part of its business, was dwarfed by the growing and investment-hungry voice business, whose switching infrastructure was in need of expensive updating.

The transition to a public corporation gave BPO-T staff more freedom to innovate [Kelly, 2010]. They worked closely with NPL and others, whilst maintaining an independent approach and supported UK academia with transmission facilities allowing connection to overseas networks [Kelly, 2011a].

Scantlebury, one of Davies' team at NPL, was seconded to the BPO-T in 1969, participating in a data communications study. He worked closely with M. A. Smith supervising four data communications-related research contracts. This research fed into a study of an 18-node hypothetical model network in 1970-71 [Williams, 1971; Kirstein, 1974], which recommended a network capable of operating in circuit-and packet switched modes [Williams, 1971]. Gardner of the BPO-T was seconded to NPL, working on the development of a Mark II network.

Like most communications providers and in-line with CCITT (Comité Consultatif International Téléphonique et Télégraphique – the international body responsible for telecommunications standards) thinking, the BPO-T favoured circuit switched data networks over packet switching but hedged their bets on both technologies. Their 1974, tender for a circuit switched network, yielded responses that were not economically viable and circuit switching development was delayed pending developments, known as System-X, to the switched telephone network. As a result, more impetus was given to consideration of packet switching [Kelly, 2010].

By 1970 the BPO-T started to specify a packet-switched network and in August 1973 an experimental packet network, known as EPSS, was approved [Kelly, 2011a; Kelly, 2011b; Broomfield, 1975]. The UK was the first country to announce a public packet network [Roberts, 1978].

M.A. Smith's team of BPO-T engineers designed a packet-switching protocol for EPSS from basic principles. The contract for development of the hardware and software for the service was placed with Ferranti in August 1973 [Feldman, 1978]: the deal covered 13 Argus 700E computers (shown in figure 3); each with 48 Kbytes of memory, to be housed in three sites, and was worth £750,000 [Aldous, 2013]. The use of a commercial processor allowed an initial experimental capability to be available by the autumn of 1975, reducing the development work required [Hadley and Medcraft, 1974; Marlow-Mann, 1979].

The trial formally started in April 1977, using three Packet Switching Exchanges (PSE) located in Glasgow, Manchester and London, with each PSE consisting of several fully-interconnected packet-switching units (PSUs) and a monitor and control point (MCP) for network control. The PSUs and MCPs were based on the Argus 700E, but the packet line cards for the customer interface were a bespoke development [Marlow-Mann, 1979].

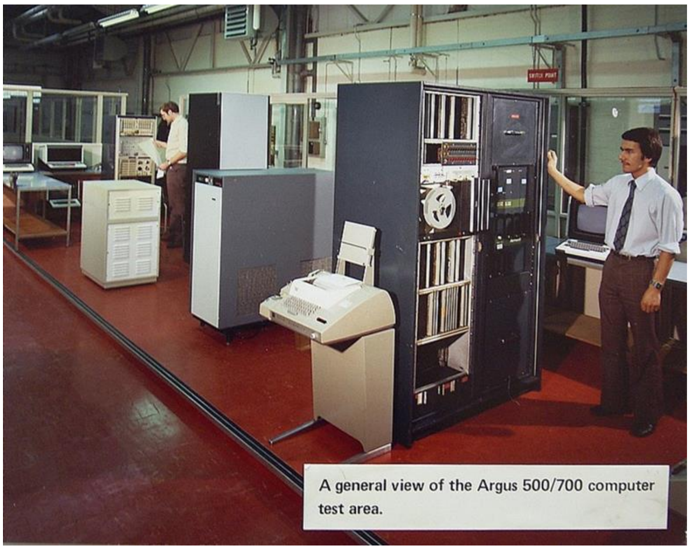

Figure 3 – photograph of a Ferranti Argus implementation (Leo Capaldi - Own work, Public Domain, https://commons.wikimedia.org/w/index.php?curid=5946200)

Users were not billed for data usage, but tariffs were devised and dummy bills issued [Kelly, 2011a]; the BPO reserving the right to charge in line with such tariffs. Prospective charges were higher than expected, potentially deterring customers from taking the service [Pelkey, 1988]. Early experience was obtained in 1974 from collaboration between NPL and the Computer Aided Design Centre, Cambridge, using EPSS protocols and by 1st March 1976, there were 40 trial participants [Després, 2010; Burgess and Lockwood, 1976].

As is typical with new and novel projects, Ferranti struggled to deliver the hardware and software for EPSS. The service initially had restricted hours of use and limited functionality and resilience, but scope and performance increased as Ferranti's software improved. For the initial release the average PSU failure rate was one fault every 90 minutes [Marlow-Mann, 1979]. The network was sufficiently stable to allow customers full access by 25th April 1977 and Ferranti delivered the final software in 1978, enabling the release of test hardware, which was used to allow dual-processor hot-standby operation and greater resilience at the London PSE.

Development was divided into four phases: provisional service from November 1975, the launched trial using evolving and reduced network equipment from April 1977; service using the final system software, from November 1977; and introduction of full network facilities in London from February 1978 [Marlow-Mann, 1979].

The experience with EPSS put the BPO-T in a strong position to influence the direction of the emerging international standard, X.25 [Campbell-Kelly, 1987; Feldman, 1978], which was ratified in October 1976. The progress in defining the protocol was rapid and it would ultimately replace the EPSS technology as the BPO-T's preferred packet switching solution [Kelly, 1981].

In 1976, a government inquiry examining data networks, emphasising national economic and social factors, confirmed the importance of such a network [Rutter 2005, p24; Linge, 2013]. After two years deliberating, it recommended the development of a national packet switched network, to be provided by the BPO-T and compatible with international standards [Rutter 2005, p24; Campbell-Kelly, 1987; Bright and Smith, 1975]. This fitted, with the plans that the BPO-T had announced for its Packet Switched Service PSS [Bright and Smith, 1975].

EPSS – the hardware and software

The Argus 700E processor, an example of which is shown in figure 4, was a 16 bit-word machine capable of addressing 64K words of memory and up to six input/output processors or channels. Each PSU comprised dual processor boards, four memory boards housing 16k words of ferrite core memory and peripheral cards. Each Argus configuration initially had 48K of core store (later expanded to 64K), two fast multiplexor (TFX) channels, and a processor interconnection channel (PIC). Processors could be arranged in live and standby pairs to provide resilience. The standby processor was continuously updated with system statistics and packet information but was isolated from the line equipment by the on-line processor which on failure allowed it to take over packet processing.

Figure 4 – Ferranti Argus 700 at the National Museum of Computing, photograph Nigel Linge

Three PSUs were installed in the London PSE, each with duplicated processors, for high availability, as shown in figure 5. Two PSUs were provided at Manchester and one in Glasgow, based on single processor configurations. The reliability and operational aspects of duplicated and non-duplicated systems could therefore be compared. A fourth node in Gresham St, London provided a billing capability and had a 10 Mbyte exchangeable hard disk drive, connecting to EPSS via a 2400 bit/s packet link [Hadley and Medcraft, 1974].

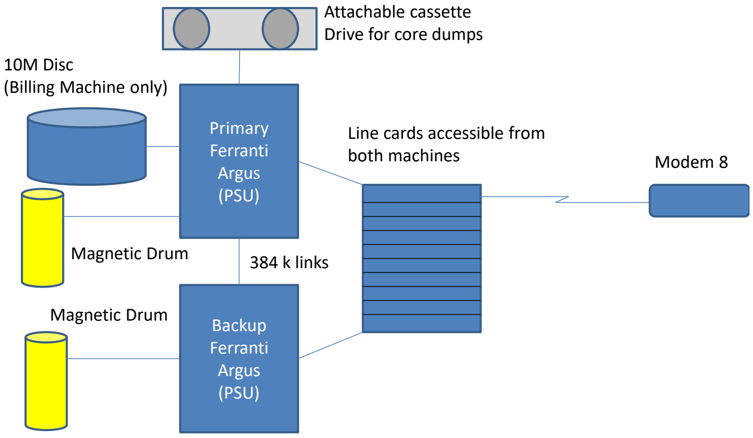

Figure 5 – Hardware Architecture of the Ferranti Argus 700E

The three PSEs were connected by 48 kbit/s trunks (three from London to Manchester, two from Manchester to Glasgow and one from London to Glasgow). The network topology is shown in figure 6. Packet mode was provided at 2400 bps, 4800 bps and 48000 bps and character mode at 110 bps or 300 bps [Feldman, 1978]. The connections available on the opening dates were as shown in table 1 [Broomfield, 1975].

Table 1 EPSS connection inventory

	London	Manchester	Glasgow
Character–telegraph 110 bps[1]	6	4	4
Direct 110 bps/300 bps	12	6	6
Dial up modem 110 bps/300 bps	51	26	26
Packet 2400 bps/4800 bps	22	15	14
Packet 48000 bps	6	2	2

Character terminals used a program-driven asynchronous character handling interface card capable of supporting speeds of 50 bps to 4800 bps. All character-line connection equipment was common and controlled by the current on-line processor.

---

[1] Hadley and Medcraft, 1974, highlight that in addition there were dial-up telex and telex services at 50 bps.

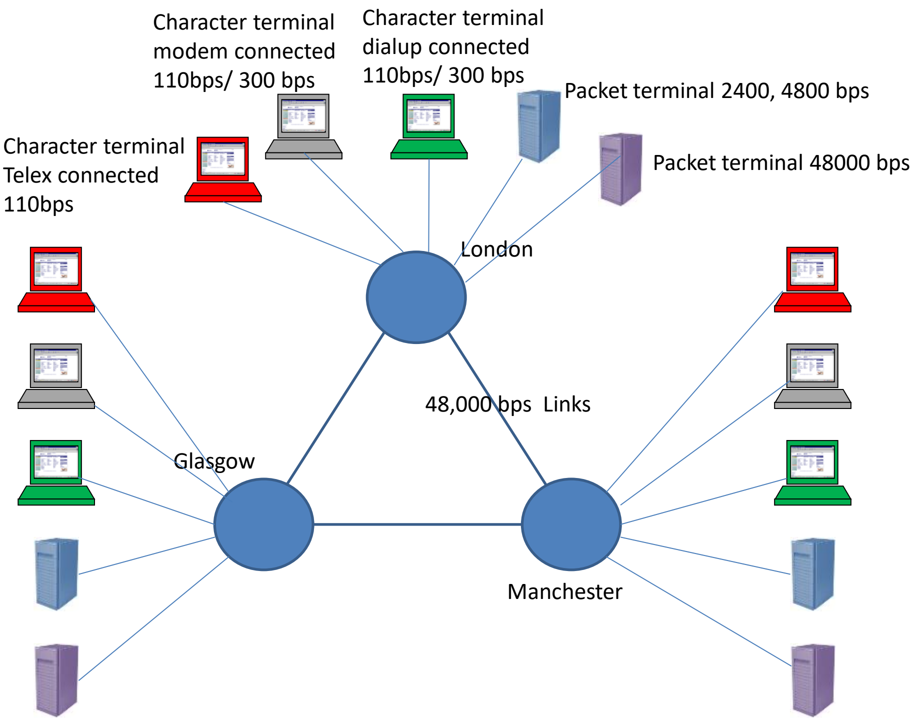

Figure 6 - EPSS topology

Packet and character lines were interfaced to the machine using line-interface cards forming a block of common interface equipment, which could be addressed by the currently on-line PSU. The line exchange capacity was shared between the PSUs, each of which addressed about one third of the available PSE line cards. Failure of one module did not result in the total loss of a service.

The packet line cards (PLC), shown in figure 7, were developed specifically for EPSS and supported circuits working full duplex synchronously at 2.4K, 4.8K, 9.6K or 48K bps for customer packet-terminal interfaces and inter-PSE trunks. The same PLC logic was used for both purposes, but wired differently according to use [Hadley and Medcraft, 1974]. The PLCs interfaced to the processor using the TFX, which was a microprogrammed module that scanned a small number of fast I/O devices and passed data between them and processor core store without processor intervention.

The links between PSUs in the same PSE or to MCPs used a PIC to provide a memory-to-memory connection. The PIC was similar to the TFX and could multiplex up to 8 full duplex lines at a throughput rate of 800 kbps [Pearson and Wilkin, 1974], enabling processors to interconnect at 384k bps. The monitor and control point (MCP) was based on a Ferranti Argus 700E with 24K of memory and had a slow printer, a visual display unit and a four-deck cassette tape system.

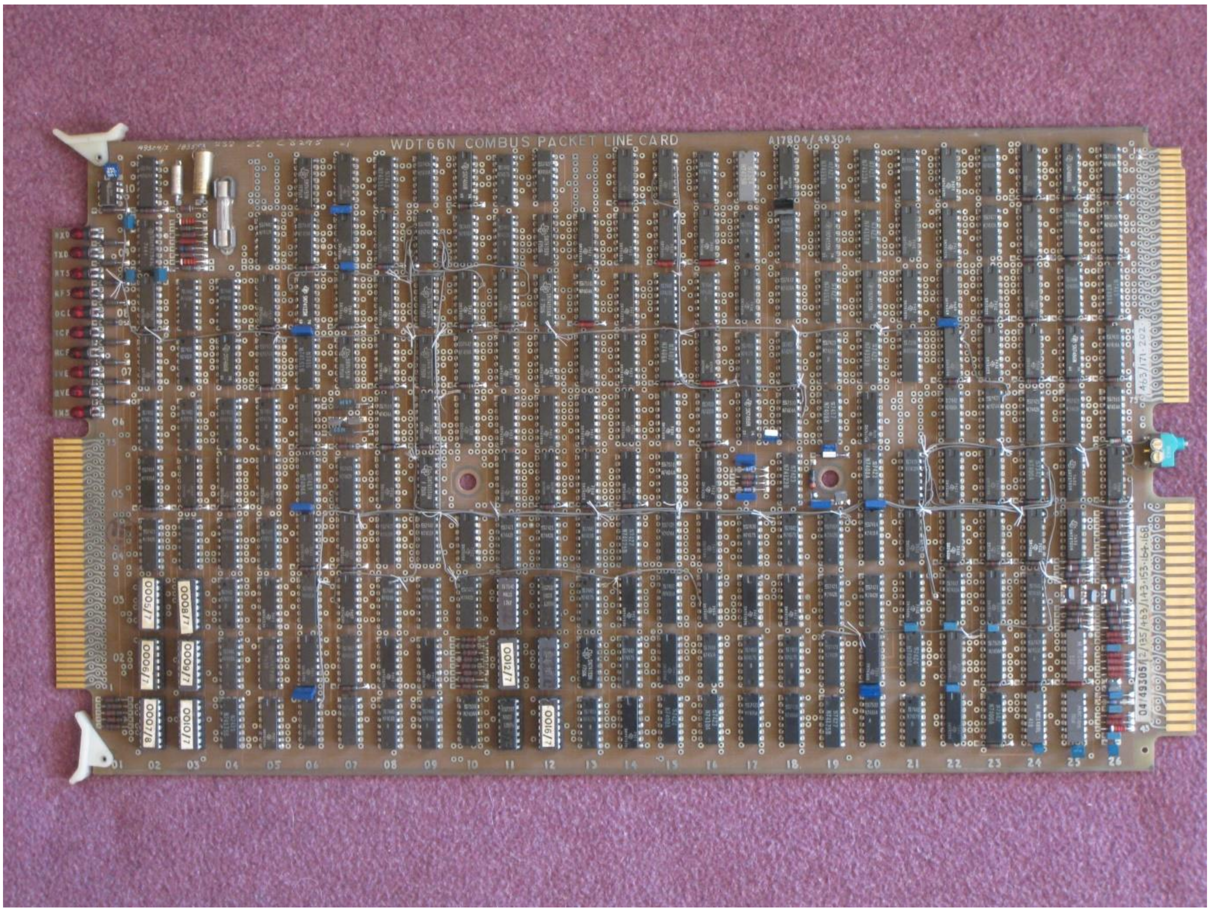

Figure 7– Packet Line Card source Chris Miller.

The EPSS software was built on the Argus 700 operating system, the layout of which is shown in figure 8 and which provided executive functions and Device Control Programs. The executive provided overall program control, interrupt handling, scheduling, memory allocation and the means of program interaction. The Device Control Programs (DCP) controlled access to peripherals [Pearson and Wilkin, 1974].

Information flowed between software modules using EPSS formatted packets and call protocols, allowing some single PSU, single drum situations to be supported by redirecting information to a slave drum at an alternative site via the network.

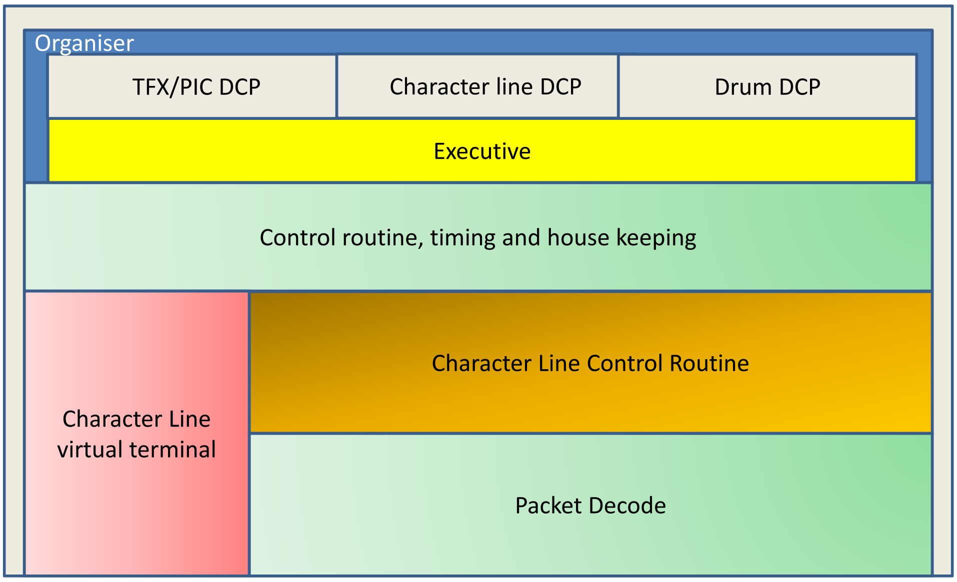

Figure 8 – Argus 700E software structure

Inter-processor communication did not use EPSS packet formats or call protocols. If a control computer failed, the standby computer was in full control within a hundred milliseconds of detecting the failure. Packets in transmission when the changeover occurs were lost and had to be retransmitted.

Programs were entered on a line-by-line basis at a terminal and editing was done using a command-line based text editor. Once a program was compiled, it was punched out to paper or Mylar tape, which was the most reliable means of storage.   This could then be read back in to the computer. Due to the operating system not implementing the code and data separation capabilities inherent in the hardware, machine crashes were not uncommon and under such circumstances the core image had to be transferred to cassette tape for analysis by specialist BPO staff.

Having described the hardware and topology of the network we now look at the protocols that EPSS used to establish connections between terminals and move data around the network.

The EPSS Protocols

The EPSS protocols were developed by the BPO-T [Kelly, 2010] and the early papers show minimal reference to work done elsewhere, which is surprising given the influence of Davies and NPL. The EPSS design was based on a nominal 256 byte packet, comprising header (including packet length), data and error-check-bit fields [Linge, 2013; Neil et al 1975a].

Call establishment was initiated by a source terminal and completed provided the call was legitimate and there were sufficient resources available to support it. Once call acceptance had been signalled, the source and destination terminals could start to exchange data.   The source PSE set up statistics and billing information collection, when it had been informed the call was accepted [Neil et al 1975a].

Once a call was established, a unique source label was allocated allowing the PSE to distinguish between simultaneous calls from a customer and permitting the called address to be omitted from subsequent packets on the customer to PSE link. The source PSE inserted the called and the calling address, in a separate field for packets travelling between PSEs. Each subsequent data packet was marked with a customer sequence number one greater than for the last packet, and cross-referred during the acknowledgement process allowing error detection [Neil et al 1975a].

Only one packet could be transmitted on the link from the PSE to the receiving terminal before an acknowledgement was received, ensuring that packets were received in the correct order. Error correction was by retransmission and packets could be simultaneously transmitted and received [Neil et al, 1975b].

The link protocol came in two forms; the standard and simplified form. The former maximised throughput, by returning a response at the earliest possible time to a packet received, interrupting a packet already in transmission if necessary [Neil et al, 1975b]. This method of data transfer was achievable because the delay characteristics of the link were measured when the link was started, and used to determine acknowledgement timing. The acknowledgement indicated whether the error was a transmission error or a packet sequencing error, or whether there were insufficient resources to handle the packet. If a PSE made three unsuccessful attempts to deliver a packet on a link, the link was considered as temporarily out of service [Neil et al, 1975b].

These techniques were not used by the simplified protocol, which delivered an acknowledgement, only if the packet was correctly received.

Having acknowledged a packet, the PSE checked the customer-to PSE link sequence number. If it was the next expected in the sequence, the packet was passed on to the next stage of validation; if not, the packet was discarded [Neil et al, 1975b].

The packet transmission procedures for inter-PSE traffic used a similar process to the simplified process, with the final exchange before delivery to the customer maintaining packet order. Up to 16 packets could be transmitted between PSEs before the acknowledgement to the first packet was received [Neil et al, 1975b]. The packet transfer procedure is summarised in figure 9.

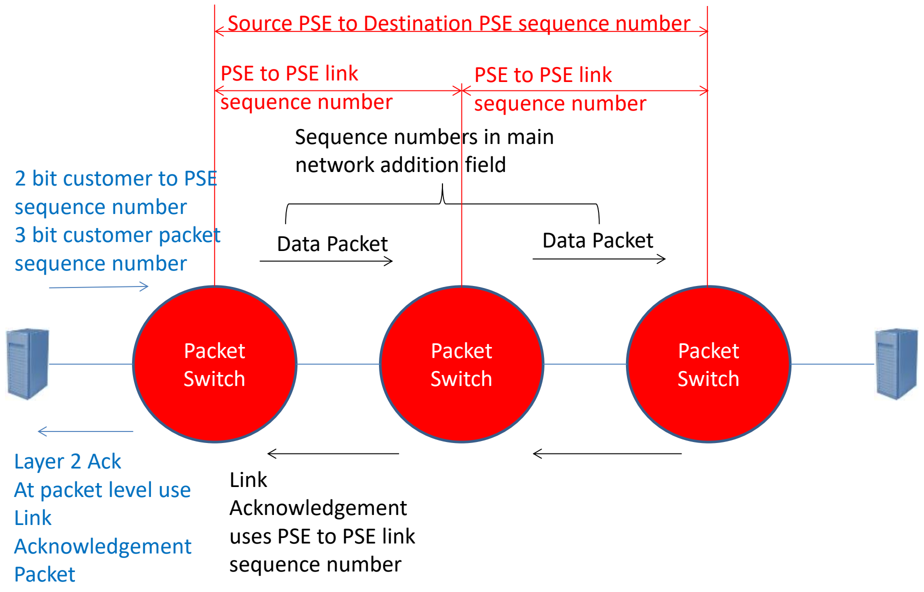

Figure 9 – The packet acknowledgement mechanism

During a call, a terminal could have a maximum of 8 packets active in the network, being able to send more as packets left at the remote PSE. Packet terminals could use the Delivery Confirmation bit in the facility field to indicate correct receipt of a packet, or group of packets. Acknowledgement was from terminal to terminal with no PSE action [Neil et al 1975a].

For asynchronous terminals, the packet assembly and disassembly facility was performed by virtual packet terminal (VPT) software in each packet switching exchange (PSE). The VPT could form all the packet types necessary for normal transmission on behalf of the asynchronous terminals. Packet headers were not transmitted to terminals by the VPT, which stripped off the header [Neil et al, 1975c].

The arrival of X.25 [32] and the provision of a European data network, Euronet [Kelly and Lee, 1978, Burgess and Lockwood, 1976] show how quickly the field was advancing and we augment the picture by considering the development of an international standard and its relationship with the approaches used on the Internet.

The move to a standard approach and competing ideas

In 1974, representatives from Britain, France, Canada, and the USA met to discuss the standardisation of packet-switched networks [Roberts, 1978], leading to, the CCITT agreeing recommendation X.25 as the definition for a packet switching protocol, in October 1976 [Feldman, 1978].

At this stage the virtual circuit model was preferred[2] over the datagram model used on the French Cyclades network, which had been operational since 1973 [Kirstein, 2009]. In the datagram model, packets are transmitted across a network independently of each other with no concept of a call, and flow control is managed by the end-user devices [Després, 2010]. The network may discard packets in the event of congestion and end stations are expected to recover from this. Its proponents argued based on the basis of cost and improved resilience, which it was thought were equally applicable to X.25.

The virtual circuit solution relieves the host of the responsibility for flow control and it was thought that the scalability of a public service based on datagrams was questionable [9]. Further, connection-based services were attractive because they allowed the specification of quality of service for each connection on a virtual circuit basis [Després, 2010].

The French liked the EPSS concept of “virtual calls”, which they used on their small Réseau à Commutation par Paquets (RCP) network in 1974 [Després, 2010], but considered the detailed EPSS protocols complex and limited. They simplified and improved on the idea, using the term “switched virtual circuit” for virtual calls. They believed that the EPSS link protocol, limited continuous transmission at full speed and could duplicate some packets on the link to the customer’s terminal; it also needed specialised hardware. The RCP design made duplication impossible, permitted continuous transmission at full speed, and was implementable using existing hardware [Després,

---

[2] The datagram approach was not well suited to the commercial aspirations of the public telephone operators, who traditionally had a usage-based charging element to their tariffs, particularly where cross-charging for international traffic was concerned. Datagram models were unfamiliar to them and would not fit well with these existing charging models.

2010]. Davies supported the view that the EPSS protocols were esoteric and held the BPO-T back [Pelkey, 1988].

Transpac, a French public network, was specified in 1974 and opened in December 1978, its initial virtual circuit specification was quickly and cost-effectively replaced by that of X.25, as the standard was very close to that originally specified. However, the Canadian implementation (Datapac), announced in July 1976, was the first public network to support fully the X.25 family of standards [Rybczynski, 2009].

Telenet were set up in 1972 by BBN with the aim of providing a commercial version of ARPANET and were led by Larry Roberts and Barry Wessler. In January 1975, after two years of effort it had raised $2m of the $25m venture capital needed to finance an American public packet-switching network [Johnson, 1976]. Telenet introduced the first public packet network service in August 1975 [Pelkey, 1988; Clarke 1982, pp 30-31]. With a few agreed minor amendments to their specification, Telenet joined the agreement on standards.

Roberts was involved with the CCITT working group working on standardising packet networks, in which both the BPO-T and the French were influential. By 1976 outline proposals were ready, but since there were still some open questions to resolve, skilful horse trading with the circuit-switching lobby was needed and X.25 was approved in 1976 [Rutter 2005, p18; Rybczynski, 2009]. This was achieved in two years, rather than the normal four-year standard's cycle and was seen as a techno-marketing-business success, allowing network operators to roll out X.25 networks [Rybczynski, 2009].

The influence of EPSS on X.25 is seen in its use of virtual circuits, the rudiments of the connection principles for asynchronous terminals, closed user groups, reset mechanisms, basic call set-up mechanism, and the use of a facilities field. Whilst the detailed specifications were very different, the BPO-T exerted substantial influence in defining X.25, playing a major role in its rapid development and adoption [Kelly, 2010; Després, 2010; Rybczynski, 2009] and also in the development of character terminal access [Kelly, 2011a; Rybczynski, 2009].

The replacement for EPSS needed to be standards-compliant and the vendor identified on short timescales. A key factor was that the BPO were mandated to consider only tenders from British firms. Four tenders were being considered, revealed by industry sources to be Plessey, Ferranti and two British Software Houses: Logica and Leasco [Marsh, 1978a].

Leaseco's offer was based on a network implemented for British Steel, which whilst not conforming to X.25 had been operational since January 1978 and could be made conformant. NPL were involved in the design of this network [Dewis, 1978]. Ferranti offered an adaption of the system used to supply the Canadian Datapac network and Logica's offer was based on the French technology proposed for Transpac and their experience in providing the Euronet network, which was however running behind schedule [Marsh, 1978b].

Plessey, using equipment sourced from Telenet, were the favourites, since the licensing agreement between the two companies allowed Plessey to incorporate Telenet hardware/software in its own solutions. In addition, a Telenet switch was already being used for the London node of an international link to the USA. The Datapac-based solution was thought to be less well tested than that of Telenet and the Transpac system had suffered a range of technical problems [Marsh, 1978b]. Ultimately Plessey won the contract, valued at £3 million and covering nine packet-switching exchanges [New Scientist 77(1090), 1978].

Having considered how CCITT protocols came to dominate, we now examine how they were displaced by the Internet.

The path to the Internet

The mobile and broadband networks that are widely used today all converge onto and are dependent on packet based infrastructures. We now consider how TCP/IP [3]displaced CCITT protocols in driving this change and illustrate this using the timeline in figure 10.

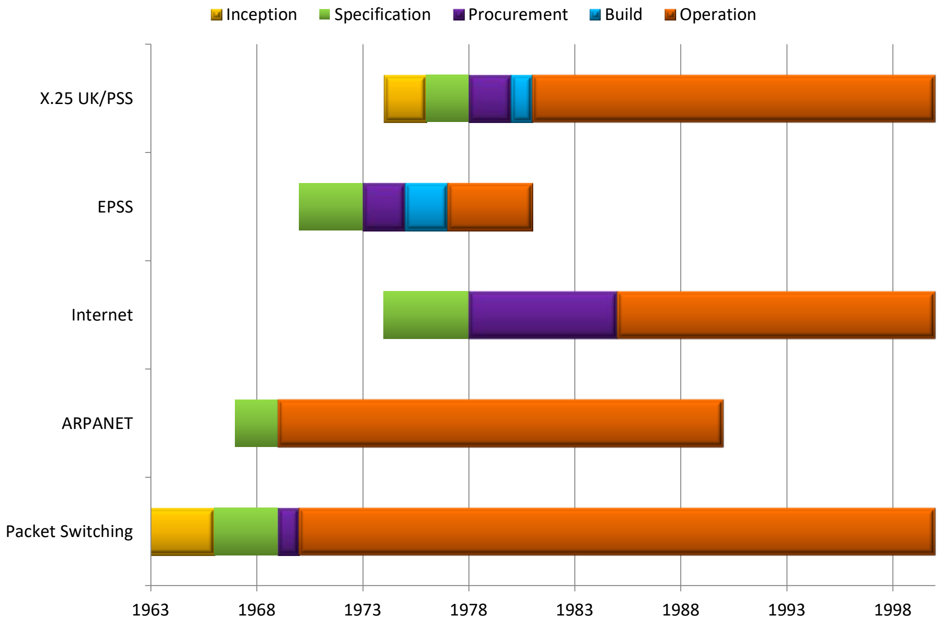

Figure 10 – timeline showing the development of packet switched technologies

---

[3] Most economists would argue that "network effects" were the most significant explanation of the rise of TCP/IP as the dominant standard. Such effects derive from Metcalfe's law which states that the utility of a service increases with the square of the number of users on that service. Clear examples would be Facebook and Twitter in the modern context, but it can be equally argued that the good old Plain Ordinary Telephone Service also flourished due to network effects. Clearly, network effects would drive the take up of packet services be they X.25 or TCP/IP, but they are a necessary but not sufficient component for achieving the market dominance. As the popularity of TCP/IP increased and its market share eclipsed that of X.25, the network effect would cause that gap to widen leading to market domination. The question is what factors put TCP/IP into a position where it had such a market advantage? This is explored in the rest of this section.

In 1972, the ARPANET used a communications protocol known as the Network Control Program (NCP) [Cerf and Kahn, 1974], which ran on the host and managed communications from the host application to the network, communicating with the network using a BBN defined end to end protocol [Carr et al, 1970; Crocker et al, 1970]. NCP was used to establish, break and switch connections and to manage control flow [Hauben and Hauben, 1998]. Whole messages were delivered from the host to the network node, where they were broken down into packets for transfer across the network links [Heart et al, 1970]. Packet delivery was assured using hop-by-hop acknowledgements and the network ensured that packets were kept in order. NCP was built on the assumption that the network was reliable [Leiner et al, 2009; Hauben and Hauben, 1998].

In addition to ARPANET, ARPA also provided a satellite network and a packet radio network and computers on all three networks needed to communicate with each other [Rutter 2005, pp 264-269]. The network could not be assumed to be reliable, requiring the transfer of the accountability for end to end delivery to the hosts. Consequently, TCP was proposed as a replacement protocol for NCP in 1974 but still required an association between the end points for the duration of a transfer, maintained by the host rather than the network. TCP was split into a lower IP layer, which provided only for addressing and forwarding individual packets (a datagram service) and TCP for flow control and recovery from lost packets, in 1978 [Leiner et al , 2009].

TCP/IP became mandated for use on the ARPANET on 1st January 1983, although it had been a defence standard since 1980 [Leiner, 2009]. The X.25 recommendation, which was further refined in 1980, in contrast had been finalised in 1976. The transition of the ARPANET to TCP/IP permitted it to be split into a defence and academic research networks. Additional research networks were developed, but the most significant was the implementation of NSFNET in 1985 [Hura and Singhal 2001, pp 749-752]. This was an academic research network using TCP/IP, linking five supercomputing centres funded by the US National Science Foundation (NSF). This core network linked to regional research and education networks that in turn connected campus networks. From 1986 it became the principal Internet backbone and the provision of interconnection points to other federal networks made it possible to close down the ARPANET in 1990.

Early commercial IP networks could interconnect to and route traffic over the NSFNET but an acceptable use policy prohibited commercial use. In 1991, three commercial network providers interconnected their networks to overcome this limitation, forming the Commercial Internet eXchange (CIX). This collaboration between providers was the beginning of the privatisation of the Internet, which was completed in April 1995, when the NSF ceased to fund the backbone, which was subsequently decommissioned [Challinor, 2000]. The funds liberated were redistributed to regional backbone networks, which were central to the expanding Internet [Rutter 2005, pp189-190]. Traffic between networks was exchanged at four Network Access Points, which NSFNET regional backbone providers had to pay to connect to. These developments provided the basis for the modern Internet.

In the 1960s and 1970s operating systems were expensive and AT&T had produced an operating system known as UNIX, which was free to academic institutions; hence many computer scientists and engineers were familiar with it. The University of California at Berkeley was encouraged to incorporate TCP/IP into its version of UNIX, enabling dispersion of the protocols to the research community, which made great use of it [Leiner et al, 2009].

Widespread development of Local Area Networks (LANs) (particularly Ethernet [Rutter 2005, p58]), PCs and workstations in the 1980s allowed the Internet to flourish. In the business world many companies needed to support proprietary protocols across the network, giving rise to a device known as the bridge router, which could handle multiple protocols, often embedding them in TCP/IP, permitting interconnection of local area networks across the wide area.

These approaches demanded higher throughput and most X.25 providers offered a service based on a top data rate of 64 K bps. Alternative protocols such as Frame Relay, offering significantly higher data rates, emerged. There was an industry perception that X.25 could never exceed 64 K bps, although there were already switches supporting X.25 at 2 M bps[Rybczynski, 2009]. A higher throughput alternative X.45 was endorsed by CCITT in 1996 [Kelly, 2011b; Rybczynski, 2009], but arrived much too late and was never commercially supported.

Personal computers had been available since the 1970s and early adopters used services such as Prestel and electronic mail. Once Internet peering services were available it was possible for several Internet Service Providers (ISPs) to develop offerings. Pipex, the first UK ISP, was set up in Cambridge and a competitor Demon was launched in 1992. Customers drove these providers to address a wider market than just the enthusiast early adopters [Rutter 2005, pp201-232; Linge, 2013]. Also networked information services like Compuserve and America OnLine had operated in the UK over dial-up connections since the 1980s [Rutter 2005, pp 201-232].

Internet growth has been accelerated by the rapid adoption of browsers and the World Wide Web, allowing easy access to information linked throughout the globe [Leiner et al, 2009]. The first browser applications were used at CERN in early 1991 and the first graphical browser supported on multiple computer platforms was available by the end of that year. Increasing use of Internet based services led to many organisations opening up consumer web access as a channel for doing retail business. Indeed Naughton [2000, p268] believes Internet evolution is driven by the forces of: the human thirst for information, the demand for entertainment and diversion and the motivation to make money from providing online goods and services.

All of these changes – summarised as a broadening of the market, expansion of an open IP infrastructure, falling technology costs, the growing consumer market, increased rate of development and the migration of standards' control away from the now privatised telecommunications companies – accelerated the migration towards IP for both public and private networks.

Historical innovation perspective

We now consider the progress made from packet switching as a concept to being the basis of a critical global infrastructure, using Kuhn's model of scientific revolution for the initial development and Winston's model to consider subsequent diffusion of the technology.

Kuhn [Kuhn, 2012] begins with the route to normal science, covering the evolution of ideas based on conventional thought. In the case of communications, following the voice-led paradigm, basing data services on a circuit switched solution would be the natural path to follow. Little normal science is designed to demonstrate novelty, it is evolutionary rather than revolutionary; circuit switching, unlike packet switching, was evolutionary. In line with Kuhn's idea of priority of paradigms, the majority of communications specialists would be educated in circuit-switched methods and were likely to be dismissive of new concepts.

Progress is driven by anomaly and the emergence of scientific discoveries; in this case the anomaly is that, unlike voice, which expects continuous use of a circuit, time-sharing computers sent data in bursts. The bursty nature of data transmission favoured the adaptation of packet-switching principles, using shorter aliquots of data to minimise queuing and optimise capacity. Khun suggests that crises arise when conditions force evaluation of an alternative paradigm and ideas are rejected only if there is an alternative. Here, the expense of circuit switching forced consideration of a different approach to satisfy the needs of the time-sharing lobby. Packet switching was a viable alternative because it was affordable, met the need and some influential engineers were familiar with it.

The new ideas tend to come from people who are not straitjacketed by the conventional thinking associated with the existing paradigm; such people are often from a different field. Baran came from a general engineering background and Davies was a computer scientist; they were not communications' specialists.

The process towards a revolution is driven by disagreements, particularly over what the important problems are. For the Post Office, universal service and an outmoded switching infrastructure, plus political changes were the major concerns. There was also a strong lobby who believed that the use of computers, enabled by telecommunications should be a major factor. Political issues do much to steer the emphasis on technology and telecommunications.

Kuhn acknowledges that it takes time for new ideas to gain acceptance and this is illustrated in the time gap between early adoption of packet switching and realisation of the information society. This progression from revolution to accepted paradigm is largely invisible to the observer and it takes time for people and organisations to let go of the old thinking. Opposition to packet switching crumbled eventually because of the increasing dependence on computers, an increasingly computer-literate workforce and public, competitive pressure from alternative providers and changes in the business and cultural environments. Once the revolution is complete progress becomes evolutionary, as the Winston analysis will show (figure 11).

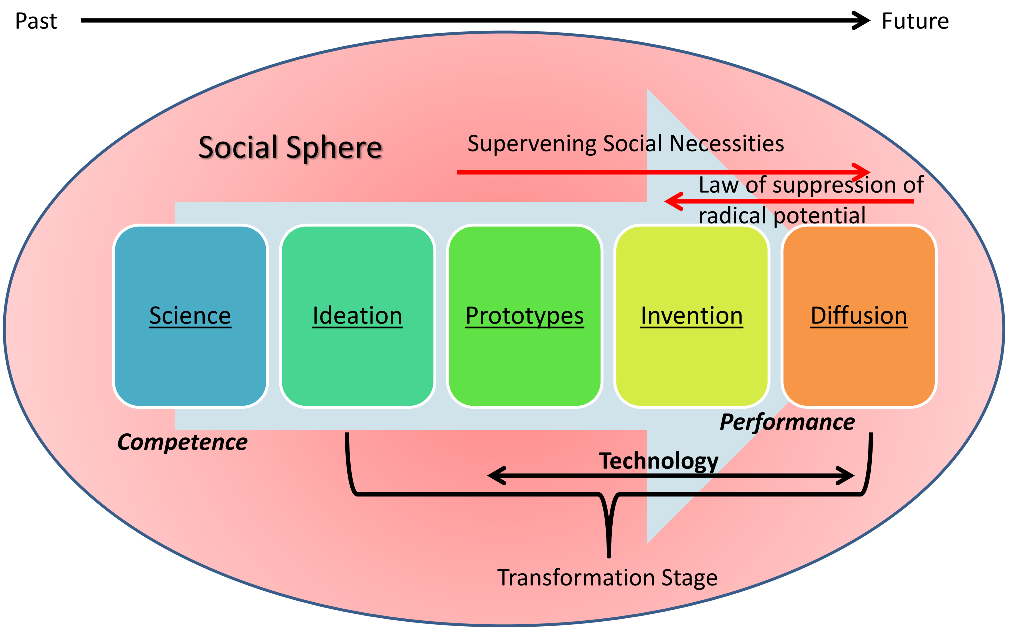

Figure 11 – Winston's model for innovation

Winston [Winston, 1998] describes a five-phase model for summarising the development of technologies. Winston was concerned how science and technology come together, mediated by social factors, to deliver new technical capability, which has major impact and is widely adopted. Success is covered by the trade-off between social necessity and the desire to remain with established technical conventions, the process of resolution of these factors leads to many innovations being lost and timescales to establish the new technology being far longer than would justify the "revolution" terms often applied to its adoption.

Whilst the early phases overlap with Kuhn's approach, Winston also covers the diffusion phase, predicated on two opposing forces, "the supervening social necessities", which drive change and "the law of suppression of radical potential", which resists it. We argue that in the early stages of packet switching, the market was limited and the "law of suppression of radical potential" dominated. However, we have highlighted services offered across the network which first appealed to the research community, then to the business community and finally to the widest possible market place. The result is the law of suppression on radical potential has been overcome and the rule attributed to the French historian Braudel that "freedom becomes value when it changes the limits of the possible in the structures of everyday life" [Keen and Mackintosh, 2001] has been realised.

In common with Winston, Fernand Braudel also saw the advance of technology oscillating between forces that accelerated it and those slowing it down. He further identified the dynamic between everyday living and the capitalist market. He concluded that that people engage with new technology when the technology of the day ceases to enable them to pursue their goals and if a new technology allows them to achieve what was formerly believed to be impossible; it will be adopted and used successfully by members of that society.

These ideas are represented in figure 12, showing that as technology became cheaper, the network offered the route to an increasing number of services, thus increasing demand through expanding the scope of the supervening social necessities. At the same time, factors that drove the law of suppression of radical potential, such as hostility to the packet-switching concept, cost of communications, limited demand for information services, the investment requirements of traditional voice telephony, the cost of computer hardware and the public's lack of expertise in computer technology have all been eroded. The diagram has limited the number of steps and some important milestones are missing, for example, the personal computer revolution. It does however demonstrate the evolution towards the information society we know today, with packet switching forming the important first step.

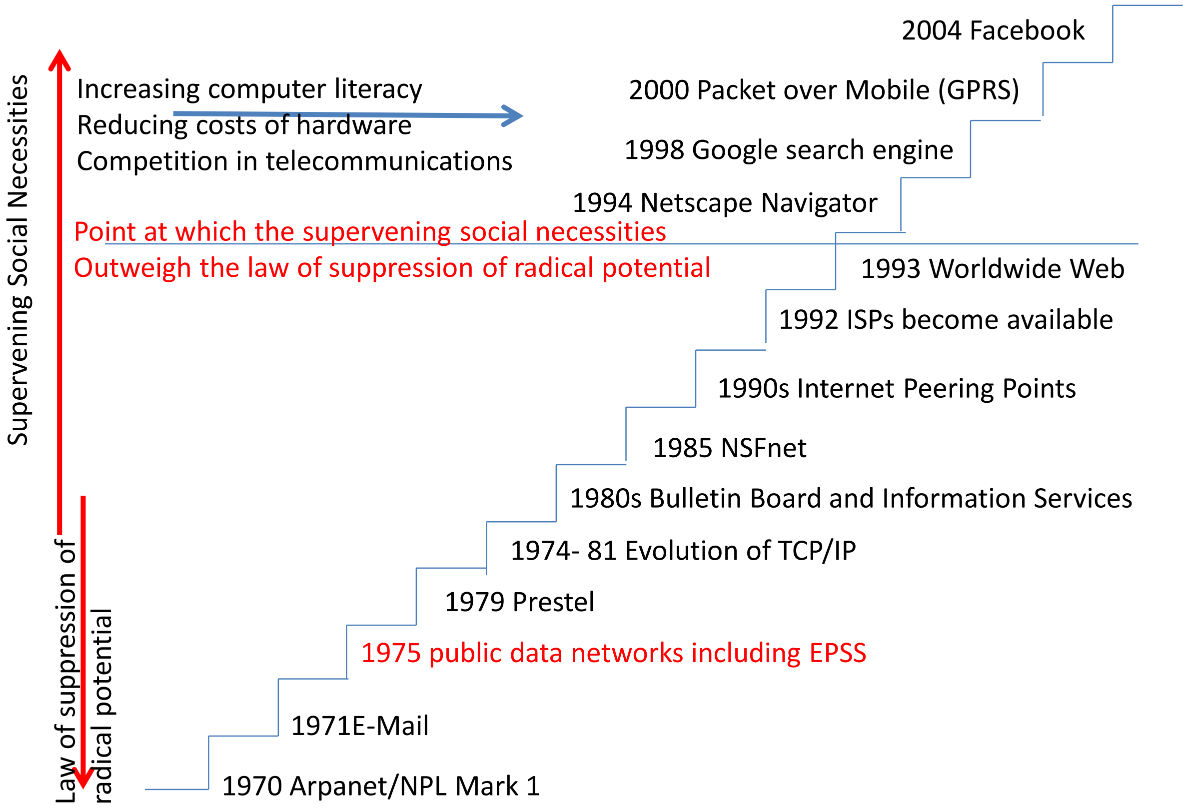

Figure 12 – The incremental innovation leading to the information society.

Conclusion

This paper has examined the development of the UK's first commercial packet-switched service, which was a very early ancestor in the evolutionary path leading to today's information rich connected world.

The evolution of these services was rapid, from a variety of national initiatives in the early 1970s to a working international standard, supported by the world's principal public network providers by 1976. The development of packet switching was achieved despite a lack of interest in the approach from telecommunications' authorities and standards bodies. We have assessed these initial barriers to development using Kuhn's model for scientific revolution.

Whilst EPSS was a key enabler in establishing and evolving the technique, it was very quickly overtaken by the standards-based option X.25, which would in the mid-1990s be displaced by TCP/IP. Yet many of the ideas trialled on EPSS later found their way into the standards. Like the early Internet and other prototype services, EPSS was based on adapted general purpose computers running specialist software. The minicomputers used were displaced in later networks by specialist microprocessor-based switches that made use of customised integrated circuits to provide high performance. Many of the concerns addressed by EPSS remain pertinent to today's networks, including providing reliable data transport over a network designed to provide a resilient service. Many of the methods used on the Ferranti Argus system are used in modern packet switching services. The big difference is that the market, at least for consumer access, has moved from closed networks, with inherent security capabilities run by a single network authority, to a situation where the network is open, available to anyone with the necessary apparatus, run by multiple providers and where security has to be added, often using standards based mechanisms.

We have traced how the approach originating with ARPA and the rise of the personal computer combined to take the initiative away from global standards bodies such as ISO and the CCITT and into the camp of the Internet standards bodies that had grown out of the ARPA organisation. Consequently, TCP/IP became the standard approach to packet switching. Commercial organisations, on the whole still tended to buy their communications services from traditional technology partners, whereas the general public initially engaged with ISPs, but now tend to buy these services as part of a "Quad play" offering from their terrestrial or mobile communications provider. Winston's model, of the Law of Suppression of Radical Potential and the driving force of Supervening Social Necessities, is particularly beneficial in understanding these points.

Despite the limitations of the equipment available and the time taken to get administrations (and the BPO in particular) interested in the technique, EPSS made a major contribution to establishing packet switching as a viable technique. It represents a vital early step on the path to the information society of today.

References

1. Abbate, J., Cold war and white heat: the origins and meanings of packet switching, Chapter 25, The Social Shaping of Technology, ed. MacKenzie, D. and Wajcman, J., 1999, second edition, Open University Press, Buckingham
2. Aldous, B., 40 Years ago .... From the Pages of Computer Weekly, Computer RESURRECTION, The Bulletin of the Computer Conservation Society, ISSN 0958-7403, Number 63 , Autumn 2013, https://www.computerconservationsociety.org/resurrection/res63.htm
3. Bright, R.D. and Smith, M.A., Experimental Packet Switching Project of the UK Post Office, NATO Advanced Study Institute on Computer Communication Networks, University of Sussex, 1975, pp 435-44
4. Broomfield, C.F., Packet Switching – The Experimental Packet Switched Service, ACM SIGCOMM Computer Communication Review, July 1975, 5(3), pp 2-11
5. Burgess, M.J. and Lockwood, J.A., Experimental Packet-Switched Service: The Customers and their Packet Terminals. The Post Office Electrical Engineering Journal, July 1976, Vol. 69, pp 122-126
6. Campbell-Kelly, M. , Data Communications at the National Physical Laboratory (1965-1975), Annals of the History of Computing, 1987, 9 (3/4), pp 221-247

7. Carr, C.S., Crocker, S.D. and Cerf, V.G., HOST-HOST communication protocol in the ARPA network. Proc. AFIPS, 1970 SJCC, Vol. 36, AFIPS Press, Montvale, N.J., pp 589-597
8. Cerf, G.C. and R.E. Kahn, A Protocol for Packet Network Intercommunication, IEEE Transactions on Communications, 1974, 22.5, pp 637-48.
9. Challinor, S., An introduction to IP networks, BT Technology Journal, July 2000, 18(3), pp 15-22
10. Clarke, P. N., Packet and Circuit-Switched Data Networks, Ph.D. Thesis, 1982, Imperial College, London.
11. Computer Networks fetch a packet, New Scientist, 16 Feb 1978, Vol. 77, No. 1090, p 436
12. Crocker, S., Carr , S. and Cerf, V., New HOST-HOST Protocol, RFC33, 12 February 1970
13. Davies, D.W., An historical study of the beginnings of packet switching, The Computer Journal, 2001, 44(3), pp 152-162.
14. Després, D., X.25 Virtual Circuits - TRANSPAC in France - Pre-Internet Data Networking, IEEE Communications Magazine, 2010, 48 (11), pp 40-46
15. Dewis, I.G., The British Steel Corporation Network, Proceedings of the NATO Advanced Study Institute held at Bonas, France, August 28 - September 8, 1978, Ed K.G. Beauchamp, 1979, Springer, Berlin, pp 345-352
16. Edgerton, D.G., Science and the Nation: towards histories of twentieth century Britain, Inaugural lecture as Hans Reising Professor, Imperial College, London, 15th October 2002
17. Fetching a Packet from the US, New Scientist, 30 Nov 1978, Vol. 80, No. 1131, p 671
18. Feldman, R., EPSS Development, Proceedings of the NATO Advanced Study Institute held at Bonas, France, August 28 - September 8, 1978, edited by Beauchamp, K.G., Springer, pp 151-155
19. Gardner, A.J. and Sandum, K.N., Experimental Packet-Switched Service: Routing of Packets, The Post Office Electrical Engineering Journal, Jan 1976, Vol. 68, pp 235-239
20. Hadley, D.E. and Medcraft, D.W.F., Experimental Packet-Switched Data-Transmission Service: Network Design and Implementation, July 1974, The Post Office Electrical Engineering Journal, Vol. 67, pp 88-94
21. Hauben. M. and Hauben, R. , Behind the Net: The Untold History of the ARPANET and Computer Science (Chapter 7), 1998, https://journals.uic.edu/ojs/index.php/fm/article/view/612/533
22. Heart, F.E., Kahn, R.E. Ornstein, S.M., Crowther, W.R. and Walden, D.C. The interface message processor for the ARPA computer network, Proc. AFIPS 1970 SJCC, Vol. 36, AFIPS Press, Montvale, N.J., pp 551-567
23. Hura, G. S. and Singhal, M., Data and computer communications: networking and internetworking, 2001, CRC Press, Indianapolis.
24. Johnson, T., Electronic post for switching data, New Scientist, 13 May 1976, Vol. 70, No. 1000, pp 351-353
25. Keen, P.G.W and Mackintosh,R., The freedom Economy, 2001, McGraw-Hill, Berkeley, pp 30-33
26. Kelly, P., Packet-Switched Data Communications Networks, The Post Office Electrical Engineering Journal, 1981, 74 (3), 216–19
27. Kelly, P.T.F., Computer Communications the early days 1966-1986, part 1, Journal of The Institute Of Telecommunications Professionals, 2010, 4(4), pp 30-39
28. Kelly, P.T.F., Computer Communications the early days 1966-1986, part 2, Journal of The Institute Of Telecommunications Professionals, 2011, 5(1), pp 54-61
29. Kelly, P.T.F., Computer Communications the early days 1966-1986, part 3, Journal of The Institute of Telecommunications Professionals, 2011, 5(2), pp 50-57.
30. Kelly, P.T.F. and Lee, E.J.B., The Telecommunication Network for Euronet, The Post Office Electrical Engineering Journal, 1978, Vol. 70, Part 4, pp. 208-216

31. Kirstein, P.T., Public Data Networks - The Next Decade, Proceedings of a IIASA Conference on Computer Communications Networks, 1974, pp 1-24
32. Kirstein, P. T., "The early history of packet switching in the UK [History of Communications]," in IEEE Communications Magazine, 2009, 47(2), pp. 18-26.
33. Kuhn, T.S., The Structure of Scientific Revolutions, 50th Anniversary Edition, 2012, University of Chicago Press, Chicago.
34. Leiner, B. M., Cerf, V. G., Clark D. D., Kahn, R.E., Klienrock, L., Lynch, D. C., Postel, J., Roberts, L.G. and Wolff, S., A Brief History of the Internet, ACM SIGCOMM Computer Communication Review, 2009, 39(5), pp 22-31
35. Linge, N., The Archaeology of Communications' Digital Age, Industrial Archaeology Review, 2013, 35(1), pp 45-64
36. Marlow-Mann, P.D.G., Operational Experience and Evaluation of the British Post Office Experimental Packet-Switched Service. The Post Office Electrical Engineering Journal, Vol. 72, April 1979, pp 43-49
37. Marsh, P., Europe switches onto data networks, New Scientist, Vol. 80, No. 1126, 26 Oct 1978, p 281
38. Marsh, P. Post Office prepares to switch on data transmission, New Scientist, Vol. 79, No. 1115, 10 Aug 1978, p 387
39. Naughton, J., A Brief History of the Future A Brief History of the Future—The Origins of the Internet, 2000, Phoenix (London)
40. Neil, W., Spooner, M.J. and Wilson, E.J., Experimental Packet-Switched Service: Procedures and Protocols; Part 1 - Packet Formats, Facilities and Switching, The Post Office Electrical Engineering Journal, 1975, Vol. 67 , pp 232-239
41. Neil, W., Spooner, M.J. and Wilson, E.J., Experimental Packet-Switched Service: Procedures and Protocols; Part 2 - Transmission Procedures, The Post Office Electrical Engineering Journal, 1975, Vol. 68, pp 22-28
42. Neil, W., Spooner, M.J. and Wilson, E.J., Experimental Packet-Switched Service: Procedures and Protocols; Part 3 - Operation of Asynchronous Terminals, The Post Office Electrical Engineering Journal, 1975, Vol. 68, pp 110-117
43. Pearson, D J and Wilkin, D., Some Design Aspects of a Public Packet Switched Network, Second International Conference on Computer Communication, Stockholm, August 1974, paper T 2244, pp 199-213
44. Pelkey, J.L., Interview of Donald Davies, Recorded May 27, 1988 London, England at http://archive.computerhistory.org/resources/access/text/2017/11/102738594-05-01-acc.pdf
45. Roberts, L., The Evolution of Packet Switching, Proceedings of the IEEE, 1978, Vol. 66, No. 11, pp 1307-1313
46. Roberts, L.G., The Arpanet and computer networks, HPW '86: Proceedings of the ACM Conference on The history of personal workstations, January 1986, pp 51–58
47. Rutter, D. J., From Diversity to Convergence: British Computer Networks and the Internet, 1970-1995, Ph.D. Thesis, 2005, University of Warwick.
48. Rybczynski, A., Commercialization Of Packet Switching (1975–1985): A Canadian Perspective, IEEE Communications Magazine, December 2009, pp 27-31
49. Williams, M.B., Developments in Data Communications, The Post Office Electrical Engineers Journal, 1971, 64, pp 70–80
50. Winston, B., Media Technology and Society: A History: From the Telegraph to the Internet, 1998, Routledge, New York

Glossary

ARPANET	Advanced Research Projects Agency (ARPA) network. ARPA is part of the American Department of Defense
BBN	Bolt, Beranek and Newman Inc - a consultancy engaged to develop and build ARPANET.
BPO-T	British Post Office (Telecommunications)
Bps	Bits per second
Carter Committee	Committee who reported on the post and telecommunications business to the UK government in 1977
CCITT	Comité Consultatif International Téléphonique et Télégraphique – the international body responsible for telecommunications standards
Circuit Switched network	A network where an association makes exclusive use of a portion of network capacity for the duration of the association or call.
CIX	Commercial Internet eXchange
DCP	Device Control Program
EPSS	Experimental Packet Switched Service
K bps	Kilobits per second.
Ferranti	UK electronics and defence company with a strong tradition in computer development.
ISP	Internet Service Provider
M bps	Megabits per second.
MCP	Monitor and Control point
NCP	Network Control Program
NPL	National Physical Laboratory
NSFNET	Network funded by the US National Science Foundation (NSF).
Postmaster General	Until 1969 UK government politician responsible for post and telecommunications.
PIC	Processor Interconnection Channel
PLC	Packet Line Card
PSE	Packet Switch Exchange
PSU	Packet Switching Unit
RAND	A not for profit policy think tank based in the United States
RCP	Réseau à Commutation par Paquets, this was an early French network.
System X	Digital switching platform deployed by BT starting in the 1980s , predominantly for voice traffic.
TCP/IP	Transmission Control Protocol/ Internet Protocol
TFX	Fast multiplexor channel
VPT	Virtual Packet Terminal
UCL	University College London
X.25	CCITT recommendation addressing packet switching procedures.

Dramatis Personae

Character	Role
Alan Gardiner	BT engineer seconded to NPL in the 1970s.
Barry Wessler	A senior manager and packet evangelist within Telenet.
Donald Davies	A researcher at the National Physical Laboratory, who came up with the idea of packet switching and was instrumental in its wider adoption.
Larry Roberts	Worked for ARPA leading the implementation of ARPANET, instrumental in the adoption of packet technology and early CEO of Telenet.
M.A.Smith	Post Office engineer who drove specification of the EPSS protocols
Paul Baran	RAND engineer one of the two engineers to develop the idea of packet switching.
Peter Kerstein	UK champion for adoption of the Internet, academic at UCL
Phil Kelly	Manager within BT who was an early champion of packet switching and instrumental in developing the X.25 standard
Rémi Després	French researcher instrumental in the development of the X.25 standards and packet switching.
Roger Scantelbury	NPL engineer who worked on NPL Mark 1 and presented its architecture at Gatlinburg
Stanley Gill	Professor of computer science and computing at Imperial College and a founder member of the Real Time Club
Tony Rybczynski	Canadian researcher instrumental in the development of packet switching and of the X.25 standards
Vincent Cerf	A US academic and Internet pioneer who was one of the principal architects of TCP/IP

Authors

Dr. Ed Smith, FBCS, FITP, University of the Third Age; Mr Chris Miller BSc.; Prof Jim Norton OBE, FREng, University of Sheffield

Dr Ed Smith, the lead author, may be contacted at edward.a.smith@btinternet.com