SYSTEM X
The history of the development of the UK's first computer-controlled telephone switching system

BY MALCOLM HAMER

Introduction

The aim of this paper is to record the events that took place in the telecommunications industry in the UK between the end of World War II and the early 1980s, as the telecommunications arm of the British Post Office (now privatized as British Telecom) set about modernizing the UK's telephone network; and to try to explain how such large amounts of money came to be spent on one part of that modernization effort – the development of System X.

In the mid-1970s, when the System X project seemed to some of the engineers working on it to be heading for disaster, those engineers often drew parallels between System X and the much-criticized Concorde project (criticized, for example, in *The Concorde Fiasco* by Andrew Wilson, 1973). But, in retrospect, it can be seen that there were big differences between the two projects. Both projects suffered from big timetable and budget overruns and were criticized as a waste of public money. However, the engineers working on Concorde in the UK and France, although they may have been frustrated by a complicated and poorly organized management structure, were trying to solve a new and extremely difficult engineering problem, that is, determining how to build a commercial supersonic airplane. By contrast, the engineers working on System X were simply trying to upgrade the components of a component subsystem (a telephone exchange) of a much larger system (the national telephone network) whose operation was well understood system. Also, they were doing so using already-proven electronic components and design principles and emulating already-successful efforts in other countries (notably at AT&T in the USA).

Not only were the System X engineers not building something completely novel, but they were also better placed to introduce new, digital components into the design of telephone exchanges than their counterparts had been at the start of their similar efforts in other countries such as the USA, Sweden, and Japan. This was because, in digital switching, the UK had roughly an eight-year lead over the rest of the world. Also, although it was not public knowledge at the time, the person who led the work that preceded the formal launch of the System X project, Tommy Flowers, was the creator of the first programmable digital computer – the secret Colossus, used in the breaking of the German Enigma code during World War II. Furthermore, the engineers at the British Post Office could have consulted (but, as far as I know, did not consult) Maurice Wilkes and his pioneering computer team at Cambridge University on the topic of software development. Wilkes and his team were the first to recognize, in 1949, some of the critical problems in writing software. They recorded these problems in the first book ever written about software – *The Preparation of Programs for an Electronic Digital Computer,* Wilkes, Wheeler, and Gill, 1951. Thus, the British Post Office team started their work with access to the most advanced switching technology and the greatest computer expertise.

The story of System X is therefore an illustration of how the leaders in a field can fritter away their lead by failing to organize themselves properly and by failing to recognize the areas in which they lacked expertise and needed to seek advice.

The outcome of the System X project

The first true System X exchange (with real telephone subscribers connected to it) was placed into service in 1981. This was almost a decade later than the dates ("early 1970s") originally talked about in 1968 when the Senior Management of the telecommunications arm of the British Post Office (BPO) approved the funding of the development of a new generation of switching systems. In fact, work on System X did not really get into full swing until 1976, as described later. Even allowing for this eight-year run-up to the start of the full development effort, System X took much longer to complete than it should have done, given the ground work that had been done in those eight years.

The two main reasons for the long time taken to complete the development of System X were:

- The failure of senior members of the BPO's System X project team to set up an effective management structure for the outsourced development activities performed by the three British manufacturers GEC, Plessey, and STC (a subsidiary of ITT); and

- The failure of senior members of the BPO's System X project team to understand and manage the major software development effort that was a critical component of developing System X.

The importance of the second reason – the mismanagement of the software development aspect of System X – cannot be overstated. The hardware development part of System X went no worse and no better than previous outsourced hardware development efforts led by the BPO (that is, the work of the various groups at the three manufacturers was poorly coordinated but the desired end result was eventually achieved). However, slow progress on hardware development was not the reason for System X being completed so far behind schedule. In fact, the constantly slipping timetable for completing the software allowed the hardware design engineers to keep re-doing their design, taking advantage of newer chips (integrated circuits) that became available as time passed. This meant that System X's hardware was more or less state-of-the-art when System X finally went into production. It was the mismanagement of the software development tasks that set the pace of the project and was a dominant factor in the enormous cost overruns.

The costs of the development of System X were almost entirely paid for by the BPO/British Telecom from ongoing revenues for telephone services. These costs included (a) payments made to the manufacturers to cover the cost of the work that they did in developing System X hardware and software from 1976 onwards; and (b) the cost of maintaining a sizable Headquarters staff to oversee the work of the manufacturers throughout this period.

It is not unusual (in terms of practices in the last fifty years in most industrialized nations) for telephone network operators that have external equipment suppliers to bear *some* of the cost of developing a new generation of switching systems. This is because the cost of the development work is so large that the manufacturers may find it difficult to raise the capital necessary to take the work to completion. Also, the manufacturers' shareholders will want to see that the network operator is committed to buying the final product. By partially funding the work, the network operator addresses these issues. (Of course, where the network operator is in itself also a manufacturer – like the old AT&T – the funding is provided entirely from network-operation revenues.)

Although partial funding of manufacturers' development work by network operators is common, it is unusual for network operators to bear *all* of the cost of developing a new generation of switching systems. After all, the manufactures will make a profit in selling the systems to the network operator once the research and development phase is completed; and they may also be able to derive further profits from selling the system to telephone network operators in other countries. Not only is it therefore unfair for the network operator to bear all the developments costs, but it is also tends to slow down the effort – because the manufacturer then has no motivation to complete the development in the minimum time possible.

Yet, in the case of System X, this was exactly what happened: the BPO (later British Telecom) bore all the costs. In fact, not only were the manufacturers reimbursed for all the costs that they incurred during the development work, the development work *in itself* became a profit-making exercise for the manufacturers. This was because, instead of charging the BPO/British Telecom the actual costs of the work (salaries, materials, office space, and a reasonable surcharge to cover general administrative overheads), the manufacturers billed on a "cost plus" basis, with truly staggering markups – in the region of 300% on top of actual salary costs. (The manufacturers were expert in extracting large amounts of money for development work from Government organizations. This was because of the many Government contracts that they were awarded, mainly for military systems.) So, with the development process itself a source of profit, the manufacturers did everything they could to maximize the amount of money that they charged for the System X work. This included not only charging inflated rates for the man-hours consumed by the work, but also hiring excessive numbers of people to do the work, and generally taking as much time as possible.

The total cost of developing System X was determined to be in excess of one billion pounds (in 1980 pounds) by 1980, putting it on a par with Concorde as a hugely expensive development funded by the British public and bringing little real economic benefit to the UK.

Although System X did eventually go into production in 1981, the rate of deployment was initially slow and from 1986 onwards it was being installed in parallel with the Ericsson AXE10 system. This outcome underscores the fact that the BPO *could* have chosen to buy the AXE10 from Ericsson in 1976, canceling System X before the one-billion-pound investment was embarked upon. (A second source of supply could have been found, such as Lucent, to avoid over-dependence on one supplier.) In fact, the option of directing the BPO to cancel System X and buy from Ericsson was discussed at length within the Department of Trade and Industry between 1977 and 1980; but in the end this was determined to be politically unacceptable, because of the political clout that the UK manufacturers possessed.

It is hard to see how giving the world one more switching system to choose from (that is, System X) benefited anyone in the long run. If, at the end of the modernization program, all the telephone exchanges in the UK had been Ericsson AXE10s, instead of one-third being AXE10s and two-thirds being System X, the UK would have been better off. An earlier start to the modernization of the network would have placed the BPO in a much better position to cope with the massive surge in demand for telephone service that occurred in the 1980s. Also, the BPO would have been able to significantly reduce the price of telephone services starting in the late 1970s because it would not have needed to build a hidden "System X development tax" into everybody's telephone charges. Any negative effects on the balance of payments or unemployment resulting from going with the AXE10 would have been very much outweighed by these benefits. And any such effects could have been minimized by arranging for manufacture of AXE10 systems in the UK. In fact, once Ericsson started to supply AXE10s to the UK, they arranged for a significant part of the manufacturing of these AXE10s to take place at factories in the UK.

A brief history of automatic telephone switching systems

The first automatic telephone switching system, known as the strowger system, was invented by Almon Strowger, a Kansas City undertaker, in 1889, and patented in 1891. The first commercial strowger exchange went into service in La Porte, Indiana in 1892.

In the UK (where telephone service, based on manual exchanges, had opened in 1879), the first strowger exchange was installed at Epsom in 1912, twenty years after the first US deployment of strowger. This marked the start of the conversion of all the manual switchboards in the UK to automatic exchanges. This conversion was a long, slow process. The last manual exchange, Portree, was replaced sixty-four years later, in 1976.

After World War II, when plans for the continued modernization of the UK network were being discussed, a choice had to be made between continuing to install strowger equipment, or changing to the newer technology – crossbar – which had been chosen by a number of other countries as a replacement for strowger, notably the USA. The decision was taken in 1949 to stick with strowger.

In retrospect we can see that it would have been better for the UK if the decision had gone the other way. The crossbar system, although also electromechanical, required less routine maintenance and repair than strowger because its moving parts moved smaller distances and with less noise and wear. The 1949 decision was partially reversed fifteen years later, in 1964, when approval was given to buy a total of about 330 crossbar exchanges (5% of the total of about 6,670 automatic telephone exchanges in the UK, including trunk and local exchanges) as a stopgap measure, to meet the growing demand for telephone service while waiting for an electronic system of some sort to emerge. By not changing to crossbar in 1949, the BPO missed a big opportunity to reduce the maintenance cost component of telephone service in the 1950s and 1960s, and to place itself in a better position to cope with the big surge in demand for telephone service that occurred in the 1970s and 1980s.

In the early 1950s the total underlying cost of a long-distance call was dominated by the cost of the transmission systems; so most of the BPO's R&D money was being funneled into transmission systems. This resulted in a new generation of analog frequency division multiplexing (FDM) systems that packed many more concurrent calls onto a coaxial cable than previous transmission systems. With the deployment of these systems, the BPO was able to hold the pounds-shillings-and-pence cost of long distance calls more or less constant – a progressive reduction in prices in real terms.

In 1958 the BPO started the implementation of STD (Subscriber Trunk Dialing, that is, long-distance dialing), with the first cut-over of STD in Bristol. The Queen made the inaugural call from Bristol to the Lord Provost of Edinburgh. (In case the call did not go through, arrangements had been made to fake the call with a manually-connected path. However, the STD call did actually go through.) STD implementation continued over the next twenty-six years, and by 1984 it was possible to dial direct from any exchange to any other exchange in the UK.

Up until the mid-1960s the purchase of strowger switching equipment from the three UK manufacturers (GEC, Plessey, and STC) was done under the Telephone Exchange Equipment Bulk Supply Agreement, which was unofficially referred to as "The Ring". The Bulk Supply Agreement was signed in 1923. This cozy arrangement guaranteed each manufacturer a fixed share of the total equipment orders. There was limited price competition, an arrangement that the manufacturers liked. During the early 1960s, the arrangement had been criticized by some observers and pressure had been applied to the BPO to scrap it and introduce competitive bidding (or at least the appearance of competitive bidding) – although it is was not clear (and is still not clear) which parties were pushing for competitive bidding and what their motivation was. Certainly, they had not thought through the implications of a changes with respect to development of new generations of equipment. The agreement was formally terminated in 1967 and a revised scheme was put in place for telephone exchanges whereby individually negotiated contracts were placed for each upgrade to a telephone exchange. However, each manufacturer was given a set of territories so that, for any given exchange, only one manufacturer would be bidding on the work. This geographic assignment of equipment supply meant that each manufacturer could keep an installation workforce in each of its assigned areas, moving them from exchange to exchange as installation orders were executed.

This new arrangement for telephone exchanges seemed, to many observers, to be only a small change relative to the original Bulk Supply Agreement. With continued collusion between the manufacturers to fix prices, and the fact that each manufacturer "owned" a set of geographical areas, it had little or no impact on the prices paid by the BPO or the manufacturers' profits. Nevertheless, the three manufacturers were very upset about the change. A whole generation of managers at the three manufacturers had lived under the rules of The Ring for their entire careers. Any change, no matter how small, was a terrifying prospect. (Perhaps they feared that this first set of changes would lead to other, more-detrimental changes in the future.)

Since 1956 development of new features and subsystems for strowger exchanges had taken place collaboratively between the BPO and the three manufacturers under the guidance of the Joint Electronic Research Committee (JERC), whose members were senior engineers from the BPO and the three manufacturers. The ending of the Bulk Supply Agreement drove the manufacturers into a state of paranoia and mutual suspicion, causing JERC to fall apart and all collaborative development to cease. This was an important change that came just before the start of what became the System X development, and I will say more about it later. However, some more background needs to be described first.

The changing cost model

By the end of the 1950s it was recognized that the transmission systems cost component of telephone calls had dropped so much that the cost of a long distance call was starting to be dominated by the operating costs of the switching systems through which the call passed between its source and destination – including both the local exchanges of the calling and called parties and the trunk exchanges that interconnected the long distance transmission paths to complete the call. A large and growing part of the total operating costs of switching systems was the labor required to maintain the equipment – repairing, checking, and adjusting all the relays and selector switches. In contemplating the growing dominance of switching system costs in their operations, the BPO recognized the need to redirect R&D funding towards building a new generation of switching systems with lower capital costs and much lower maintenance manpower requirements.

As side note, it is interesting to observe that this swing towards favoring R&D on switching systems (versus R&D on transmission systems) is part of a "natural" repeating cycle, in which the investment pendulum swings between transmission systems and switching systems. Since the end of World War II, we have seen the following cycles:

- *1946 to 1960:* Concentration of R&D on transmission systems, since the cost of these dominated the cost of long-distance calls. This resulted in widespread deployment of high capacity FDM systems across the trunk network, greatly shrinking the transmission cost component of long distance calls.

- *1960 to 1981:* Concentration of R&D on switching systems, since the cost of these had started to dominate. This investment in R&D resulted in the TXE series of exchanges and, finally, in System X.

- *1981 to 1995:* Concentration of R&D on transmission systems again, resulting in widespread deployment of optical fiber systems domestically and internationally.

- *1995 to whenever this phase ends:* Reflecting the growing importance of data communication, this phase started with huge investments in IP data switches (routers), to create the present day Internet. In telephone systems, there were large investments in cellular network infrastructure, including switches and cellular base stations, which, in addition to radio-communication subsystems, contain a lot of switching functions. Plans are being discussed for the next generation of terrestrial telephone switches, although at the time of writing (2001) it is unclear whether these will be based on traditional PCM switching techniques or "voice over IP" technology.

In 1960, with the BPO prepared to pump money into R&D on switching systems for the first time since the 1930s (when the last major improvements were made to strowger), it was important that these R&D efforts be led by someone with vision, determination, and the ability to command the respect of junior technical staff. As I mentioned in the Introduction, the BPO was fortunate to have Tommy Flowers as leader of the team at the BPO's Dollis Hill Research Station who were to undertake this work.

Flowers had been a key technical person in the German military code-breaking effort at Bletchley Park in World War II. He had built Colossus, the ultimate code-breaking engine, which was the first true digital electronic computer. Colossus played a critical role in decrypting German High Command messages in the months leading up to, and during, the D-Day invasion. Unfortunately, the extreme secrecy of the Bletchley Park work meant that neither Flowers nor his colleague, Alan Turing, were publicly recognized as the inventors of the digital computer after the war. And even when information about Alan Turing's work at Bletchley Park started to emerge in the 1980s, very little was said about Flowers's contribution. Recognition of Flower's role in building the world's first computer did not come until the last few years of his life, when a number of books and television documentaries fully described, for the first time, his pivotal role.

Flowers's experiences with electromechanical code-breaking machines at Bletchley Park, and the technical success of the Colossus, had left Flowers determined to build the next generation of switching systems using nothing but electronic components – even if, in the short term, this led to higher capital and maintenance costs. He believed that, over time, improved electronic technologies would emerge to solve these problems. Some early Dollis Hill designs of switching systems used dozens of electronic components to replace a single relay. Although, in retrospect, we can see that Flowers was right about the evolution of electronic component technology, at the start of the 1960s not everybody at Dollis Hill shared his faith. Nobody, not even Flowers, could have foreseen the rapid increase in sophistication and fall in prices of integrated circuits that would result from their consumption by the computer industry over the next fifteen years. So, although we now see that Flowers's insistence on a purely electronic switching system, with no moving parts, as visionary, it is also clear why others felt it was impractical and misguided.

Flowers's first attempt to demonstrate the feasibility of an all-electronic telephone exchange was a pilot system installed at the Highgate Wood exchange in London (about 4 miles from Dollis Hill) in December 1962. This pilot resembled the Colossus code-breaking computer in that it consisted largely of a combination of vacuum tubes (radio valves) and neon tubes. Although it performed erratically in the first few weeks, it was eventually stabilized and stayed in service for several years. It was common in Dollis Hill in the late 1960s to refer to Highgate Wood as a "technological disaster", rather than a useful learning experience. This was really quite unfair (and I regret that I personally subscribed to this view at the time).

It should be noted that, up to 1962, there had been no serious talk about using a general-purpose computer to control the switching functions in a telephone exchange (stored program control). Highgate Wood used non-programmable "hard-wired" control subsystems. The design of Highgate Wood did, however, use a design approach in which "control" functions were separated from "switching" functions. This marked an important step towards stored program control.

The idea that control functions could be separated from switching functions emerged gradually in the 1940s and early 1950s. The evolution of switching systems during the first half of the 20th century had been based mainly on strowger technology. In the 1940s a number of manufacturers started to build crossbar systems as an alternative to strowger. The UK manufacturers were among these, creating crossbar products for the overseas market. In designing crossbar-based telephone exchanges, switching engineers began to see that a telephone exchange contains two conceptually separate functions – "switching" and "control". *Switching* means the connecting of one transmission path to another, for example, one subscriber's local line to another's, or a subscriber's local line to a trunk that goes to another exchange. *Control* means taking commands from the user (in the form of strings of dialed digits), interpreting these as a request for a particular service, and acting on the switching components to cause the necessary paths to be set up. Control also includes managing the communication of information to other exchanges downstream in the communication path (a process called "signaling"), and the recording of the details of calls for billing purposes. In strowger exchanges the control and switching functions are so closely intermeshed that it is hard to see them as separate functions. But in crossbar exchanges the control functions are, at least in part, implemented as separate subsystems. The hard-wired control subsystems of Highgate Wood were, in a similar way, distinct from the switching components.

After Highgate Wood, work continued on all-electronic exchanges, culminating in the cutover of the Empress pilot all-electronic digital "tandem" near Earls Court on 11 September 1968. (A "tandem" exchange is an intra-metropolitan-area exchange that switches calls between one local exchange and another. A tandem exchange, like a trunk exchange in the long distance network, has no subscriber lines connected to it. It handles only "through traffic".) The Empress tandem was a true all-electronic digital switch with hard-wired electronic control subsystems, built using a combination of integrated circuits (which had just started to become available) and discrete electronic components. *Empress was the first digital switching system in the world to go into service in a public network*. (By contrast, AT&T did not put a digital trunk exchange into service until January 1976.)

After the 1968 cutover of Empress, the next digital trunk/tandem exchange to go into service in the UK was the first System X tandem exchange at Baynard House in London *twelve years later*, on 1 July 1980.

This twelve-year period of stagnation between Empress and Baynard House is the most vivid illustration of the time that was lost through the mismanagement of the software development efforts that were required to reproduce, in software, the hard-wired control functions of Empress and Highgate Wood.

Reed-relay switches

Because of fears that the work at Dollis Hill would not create a deployable electronic telephone exchange quickly enough, a separate parallel effort was started under the umbrella of JERC in the early 1960s. This effort aimed to create a family of telephone exchanges that would require less manpower to maintain than strowger (and would also require less physical space), thus achieving the goals of the Dollis Hill work, but with less technological risk.

The result of this effort was a design for a switching system using *electronic* control subsystems (hard-wired) plus an *electromechanical* switching matrix. The switching matrix was based on reed relays. Tests had shown that reed relays were even better than crossbar in terms of maintainability because they have no adjustable parts. The contact points are enclosed in a glass envelope and are fixed in the glass during manufacture. The activating electromagnet is wrapped around the glass envelope. Several other countries had development work underway on reed relays, including AT&T in the USA and Ericsson in Sweden.

Work on the reed-relay-based systems went smoothly, and the first exchange, called a TXE2 (a small/medium sized local exchange), went into service at Ambergate, Derbyshire in 1966.

The TXE2 achieved its goals of reducing the amount of maintenance labor required to keep an exchange functioning. It was found that the TXE2 achieved a figure of 0.11 faults per subscriber line per annum. (This rate was expected to be more or less constant over a wide range of exchange sizes, including the planned larger versions of the TXE design.) This was an excellent result. It was almost 100% improvement over crossbar, which had a fault rate of 0.21, which in turn was an improvement over strowger which had a weighted-average fault rate over all classes of exchange of 0.62. (The number of faults per subscriber line per annum was 1.00 for the Director System – the variant of strowger used in London and other large cities; 0.42 for medium-sized Non-Director exchanges; and 0.28 for small rural exchanges - "UAX"s and "SAX"s.)

The better than 5-to-1 improvement in fault rate achieved over strowger meant that, in the worst case, if the work at Dollis Hill failed to produce a deployable switching system, the BPO now had an "insurance policy": deployment of the TXE2 (and its planned big brother, the TXE4) would allow the BPO to absorb a five-fold increase in the number of telephone lines in the UK (from 7 million in 1967 to 35 million) without any increase in the size of its exchange maintenance workforce – provided that the pace of deployment of TXE systems stayed *ahead* of the growth in the number of lines, so that the network progressed towards total replacement of strowger at a steady rate. (The number of 35 million lines would not be reached until the late 1990s.)

The TXE2 also achieved savings in space, since it was considerably smaller than the strowger equipment that it replaced. Interestingly, not all the potential space savings that could have been achieved in buildings built specifically to house new TXE2s were achieved, because the size and layout of toilet facilities were dictated by union agreements that could not be renegotiated in time to reflect the lower frequency of staff visits to unattended sites. (Visitors to the first TXE2 sites would remark "the loos are bigger than the equipment room!")

During the middle of the work on the TXE2, the BPO had already reversed its 1949 decision on crossbar and had, as a stopgap measure, ordered crossbar exchanges from the manufacturers (who were already supplying crossbar to overseas customers). Interestingly, this decision caused a stir at the time. In spite of the active involvement of Plessey in the TXE2, and in spite of the fact that Plessey was to be the first supplier of crossbar in the UK, staff in the part of Plessey that manufactured strowger equipment were shocked by the move. Crossbar was built in a different Plessey factory to the one where strowger was built. The strowger factory was in the constituency of Huyton, near Liverpool. This was the constituency of the then-newly-elected Prime Minister, Harold Wilson. Reportedly, a senior member of the BPO was summoned to 10 Downing Street to explain to Wilson why the BPO had decided to move away from strowger, the production of which provided a number of his constituents with their livelihood.

Fortunately, Wilson's attempts to strong-arm the BPO came too late. The first crossbar exchange went into service at Broughton, Lancashire in 1964. This was known as the TXK1, based on Plessey's export product, the Plessey 5005 crossbar system. Acting under JERC, Plessey passed its design to GEC so that they could also manufacture the TXK1. Later, in 1971, when a larger switch was required for metropolitan areas and trunk exchanges, ITT's crossbar product was deployed under the name TXK3.

So, by 1967 the BPO had two "insurance policies" in operation in case the hoped-for all-electronic exchange did not emerge in time: crossbar (the TXK switching systems) and reed-relay switching systems (TXE). Work was getting started on a large reed-relay exchange, the TXE4, for metropolitan areas and trunk exchanges.

The end of The Ring

The announcement, in 1967, of the ending of the Bulk Supply Agreement, after over 40 years of collaborative development and price-fixing, brought development activities under the JERC umbrella to a halt. JERC fell apart. This left the TXE4 development unfinished. (About five years later the BPO persuaded STC to resume work on the TXE4 and finish it off. They did; and the first TXE4 went into service at Birmingham in February 1976. STC made sure that the TXE4 was "theirs" and did not share the final design with GEC or Plessey.)

At around the same time that the Bulk Supply Agreement was ended, Tommy Flowers retired from the BPO, handing the baton to L.R.F.Harris (Roy Harris). Soon after that, Roy Harris went on extended medical leave – no details were every officially released, but it was rumored in Dollis Hill that this was due to a nervous breakdown – and did not return to work for almost two years. The switching systems research team at Dollis Hill that Harris should have assumed control of came to be organized into two groups – R8 and R15. R8 was headed by Nick Martin and R15 was headed by Charles Hughes. Martin and Hughes reported direct to a Deputy Director (H.B. Law) who had other responsibilities besides switching systems and who involved himself very little in the day to day work of R8 and R15. The separate roles of R8 and R15 were not clearly defined and there appeared to be considerable overlap between the responsibilities of the various teams within R8 and R15.

In the middle of this turbulence, the teams at Dollis Hill intensified their interest in stored program control (SPC), that is, the control of the switching components of a telephone exchange by a general purpose computer, running software written to perform the control functions. These are the control functions that had, up until then, been distributed among relays and selectors in strowger and crossbar exchanges; and, more recently, had been implemented in purpose-built electronic circuits in the Highgate Wood pilot, the Empress pilot, and the TXE2.

It is unclear how the interest in SPC at Dollis Hill became as strong as it did in the second half of the 1960s. Although Flowers was generally enthusiastic about computers (which was hardly surprising, given his work on the still-secret-at-that-time Colossus), during the 1960s Flowers himself had become increasingly skeptical about the immediate advantages of SPC. As early as 1961 he had visited Bell Labs to learn about the work there on the first SPC system, the No.1 ESS (generally shortened to "ESS1"). In a report that he wrote after this visit, he recorded his doubts about the claim of the ESS1 project leaders that SPC would give telephone exchanges vastly increased flexibility because the control software would be "easy to modify". Presumably he had spoken to the ESS1 programmers and heard about some of the serious problems that they were facing in creating the software for the ESS1. In the conclusion of his report, he urged his staff at Dollis Hill to exercise caution in trying to incorporate SPC into the design of the new generation of exchanges for the UK. However, his warnings seem to have been ignored. The pro-SPC lobby at Dollis Hill attempted to discredit Flowers in the next few years. The accepted wisdom among the more junior staff when I joined the BPO at Dollis Hill in 1970, after Flowers had retired, was that Flowers was a bit of a fuddy-duddy and that Highgate Wood had been a total disaster.

Following his retirement from the BPO, Flowers took a post-retirement position at STC as a senior technical advisor. In 1971, as a representative of STC, he presented a paper at a meeting of the Institution of Electrical Engineers, based on his work at STC. His paper was an even stronger warning about SPC than his 1961 internal BPO report on his Bell Labs visit. In his IEE paper he proposed, in effect, a blended approach to control system design, in which those control functions that would be excessively complex to reproduce in software would, instead, be implemented in hard-wired logic – working in combination with a more limited SPC control subsystem.

It is interesting that Flowers was giving these warnings at a time when STC's parent, ITT, was embarking on the development of its own SPC System, System 12, whose runaway costs and timetable would ultimately be a major factor in the collapse of the ITT empire.

At Dollis Hill, following the successful Empress pilot in 1968, work on digital switching matrices started to take a back seat. SPC became the thing that all the new graduate recruits wanted to work on when they joined the Dollis Hill team, starting with the Autumn 1968 intake of new graduates.

When I visited Dollis Hill as a student in July 1969 (one year prior to finishing university), I saw a model SPC system consisting of a reed relay switching matrix and a commercial computer (a Honeywell 516). The team demonstrated the results of their programming efforts thus far. It was possible to pick up the handset of one of the telephone instruments connected to the matrix and receive dial tone. However, in order to complete a call to one of the other telephones, it was necessary to feed a punched tape into the Honeywell to simulate the dialing of digits. The team explained that they were still working on the digit-detection and analysis programs.

At the time I thought "What is taking them so long? How difficult can it be to emulate the electronic control circuits of an exchange programmatically?" My misguided intuition about the complexity of SPC was probably fairly typical of most people at Dollis Hill at the time. They had reasoned thus:

- The relays and selectors that form the distributed "control" element of an electromechanical exchange are simple devices, whose individual actions are easy to describe and whose interactions with one another are also easy to describe.

- A program to emulate the action of any one control device is just a few lines of code.

- There are many thousands of components in a telephone exchange that collectively perform the control function. However, these are just multiple instances of a small number of distinct types of device interacting via simple electrical interfaces.

- Therefore a computer program that embodies the complete control function should not take more than a few months for a small team to write.

They failed to understand that the simple logical functions performed by individual relays and selectors among thousands of similar devices, when combined via a network of interaction paths, forms an immensely complex and powerful information processing system. They also failed to understand that a computer, although very much faster in its internal electronic circuitry than relays and selectors, would face a big challenge in handling thousands of events in parallel and in "real time".

Nobody had thought to ask a number of very important questions, to which answers might have been obtained by, for example, talking to AT&T (who by then had completed the ESS1 and had been through a number of major overhauls of its software design) – questions such as:

- How do you measure the complexity of a function such as controlling a telephone exchange?
- Using an appropriate unit of measurement, what *is* the complexity of the control function?
- How does this translate into man-months or man-years of programming effort?
- How do you organize a team of more than a few programmers?

Digital switching and digital transmission

Before continuing with the story of the work on SPC, it is worth reviewing in more detail the state of digital switching and digital transmission technologies in the late 1960s.

Although there might have been unpublished work on digital switching going on in the background at Bell Labs in the late 1960s, in terms of work taken through to a public network pilot, *the UK had a lead of almost eight years* in 1968.

The fact that AT&T had not enthusiastically pursued digital technology was largely because AT&T's accountants had determined that a network-wide conversion to digital transmission and digital switching would not become economic in the foreseeable future.

In the UK, following on from the development of high-capacity coaxial cable analog transmission systems using frequency division multiplexing (FDM), the BPO and the manufacturers turned their attention to pulse code modulation (PCM) transmission systems (that is, digital transmission systems). The concept of PCM was invented in 1937 by Alec H. Reeves, working at the Western Electric Company (part of ITT) in Paris. ITT filed a patent for PCM in France in 1938, in the UK through STC in 1939, and in the USA in 1942. However, the electronic components required to build a reliable, economic system did not exist at that time. (The patents showed examples of PCM systems built using vacuum tubes.)

The UK work on PCM led to the creation of a practical 24-channel PCM transmission system that operated over twisted pair cables and whose deployment on trunk routes in the UK started in 1964. The deployment of these PCM systems was concentrated on the shorter trunk routes within metropolitan areas and between adjacent towns. By using twisted-pair cables that were previously used to carry individual voice paths, implementation of PCM on these routes allowed the BPO to quickly and cheaply increase the capacity of these existing cables, without investing in new coaxial cables. (Later versions of PCM transmission equipment operated on coaxial cables, replacing the FDM systems.)

Unfortunately, some of the favorable economic impact of this pioneering work in PCM transmission systems was lost by poor execution when it came to production. Arrangements for the supply of transmission systems differed significantly from those for telephone exchanges. Transmission systems were not purchased in bulk and installed by BPO staff. Instead each job was put out to tender, so there was at least an element of competition. For example, the BPO would ask for bids to implement 10,000 voice channels between London and Birmingham. There were five manufacturers supplying transmission systems to the BPO: STC, GEC, Plessey, Pye (a subsidiary of Philips), and AEI. Contracts would be awarded based, at least in part, on price. (Of course some price fixing was going on among the suppliers; but the use of five suppliers made this less extreme than in the case of switching systems.) Also, the open bidding arrangement meant that manufacturers did not come to "own" parts of the BPO network in the way that GEC, Plessey, and STC "owned" the telephone exchanges in geographic areas.

Given this situation, a sensible approach for manufacture of the PCM systems would have been for the BPO to dictate the detailed equipment design down to the component level, so that parts from the different manufacturers would be interchangeable. However, sense did not prevail: the five manufacturers were allowed to create five functionally identical systems whose design, at the level of printed circuit boards, differed significantly. This meant that every BPO transmission systems repair center had to: (a) maintain five separate inventories of parts; (b) have five separate testbeds for testing repaired parts and subsystems; (c) have five sets of documentation; and (d) train their staff on five different systems. So, even though the deployment of PCM systems produced some savings for the BPO, the existence of five different designs seriously eroded the potential savings that could have been achieved with a single standard design. (Interestingly, the BPO did not appear to have considered outsourcing the repair work to the manufacturers. They preferred to stick to their traditional in-house repair approach. It is possible that the true costs of supporting all these variants were not understood until it was too late to consider outsourcing.)

Although everyone recognized that digital *transmission* systems represented a significant improvement over analog systems in terms of cost and reliability, in 1970 there were two different points of view on digital *switching* systems in the BPO:

- ***The "let's go all-digital" view.*** According to this view, AT&T had "got their sums wrong" (in determining that a network-wide conversion to digital technology would remain uneconomic for the foreseeable future). A completely digital network of digital transmission systems *and* digital switching systems would be the most economic solution for *any* national network within the next five to ten years. It followed from this view that switching systems work should focus on creating *all-digital switching systems* as quickly as possible. This was certainly the view of Flowers's few remaining disciples at Dollis Hill. Another consideration that reinforced the all-digital view was the desire to create a network that could handle more than just telephone calls – a type of network that later came to be known as an ISDN (Integrated Services Digital Network). This idea of a multi-services network was one that many of the other European telephone administrations and manufacturers were also enthusiastic about.

- ***The "hybrid switching system" view.*** According to this view, a local exchange would need to consist of a mixture of analog and digital technology, at least for the next ten years. (Trunk and tandem exchanges, however, could be all-digital.) An all-digital local exchange would not be economic because it would require analog/digital convertors on a "one-per-subscriber-line" basis. (Subscriber lines were at that time, and in most cases still are, pairs of copper wires – in underground cables, in overhead cables on poles, or "open wires" on ceramic insulators on poles.) In order to avoid having one analog/digital convertor per subscriber line, local exchanges would need to combine an analog switching matrix (for example, reed relays) with a digital switching matrix – at least until analog/digital convertors could be made much more cheaply, or until the local distribution network had itself become digital all the way to the home or office. (Another group in Dollis Hill was in fact working on digital local distribution. They saw this as likely to be a reality in the UK within 20 years – by about 1990. This was a very optimistic prediction: the complete replacement of the analog copper-wire local network by digital networks based on optical fiber cables is unlikely be complete until 2020 or later.)

These two views – total versus partial digitization of the network – remained unreconciled for the next eight years, although most of the work of R8 and R15 was based on the second ("hybrid") view. It was only in the late stages of System X development that the all-digital view became accepted by all parties.

The initial design for System X that emerged from work between 1968 and 1976 included a "DSS" (Digital Switching Subsystem) that could be used as both a trunk switch for inter-exchange traffic and as the core of larger local switches. The DSS would be surrounded by reed-relay switching modules that "concentrated" traffic from copper-pair subscriber lines onto a smaller number of trunks for which "one-per-trunk" analog/digital convertors would be economic. The smallest local switches would be nothing but reed-relay based concentrators, with one-per-trunk analog/digital convertors. The convertors would be connected to PCM transmission links that would carry the calls to the DSS in the exchange in the nearest large town. Since the local concentrator could not connect calls between two local subscribers' lines, calls between two subscribers on the same concentrator would therefore be "tromboned" through the large exchange in the nearest large town.

It was not until the last few years of System X's development that it was recognized that the analog switching part of System X could be dispensed with, and the system made totally digital, without making the capital cost of the system unacceptably high.

Almost eight years after the BPO's Empress pilot digital trunk switch went into service, the first ESS4 digital trunk switching system was placed into service by AT&T in Chicago, on 17 January, 1976. In those eight years AT&T had reversed its earlier position on digital technology and was headed towards an across-the-board conversion of its long-distance network to digital. In 1976 the System X DSS was still a paper design; so it is clear that the UK had, by then, lost its eight-year lead.

This eight-year lead had been lost not because of any lack of interest in digital technology, nor because of any problem with use of the technology in the Empress pilot. After Empress, both the BPO and the manufacturers continued to build laboratory models of digital switches. This kept their knowledge and expertise up to date. If, at any time, someone had said "build a digital switching matrix now", the UK was capable of building a switch every bit as good as that of ESS4. The problem was that the timetable for rolling out System X was now totally determined by the completion of the development of the software that would control both the DSS and the reed-relay local exchange modules; and this timetable was out of control.

A brief aside: ISDN

I mentioned earlier that the desire to build a multi-services network (that could handle data as well as voice) was one factor that the supporters of the all-digital approach regarded as part of the case for an all-digital design. Although what was later called "ISDN" did not, in fact, turn out to be a significant consideration in the final design of System X, ISDN eventually (in the late-1990s) became an important service (although not as important as was originally expected). It therefore deserves mention here. In order to avoid interrupting the main flow of the story of System X later in this paper, I will now describe how ISDN emerged and examine the extent to which the multi-services network concept actually affected early design decisions for System X.

In 1970, when the idea of a multi-services network was first discussed in Dollis Hill, the "state of the art" in data communication was dial-up through the telephone network at a speed of 300bps (or leased circuit operation at 9600bps). Dial-up connections were starting to be used to connect remote users to computer time-sharing services, typically from a teletypewriter.

There was a lot of interest in videoconferencing at Dollis Hill in 1970. In fact, this was one of the Director of Research's favorite topics. There was research going on into digitization of video signals and into compression algorithms to reduce the bandwidth required to convey these signals. However, with the computing power available at the time, nothing remotely approaching the efficiency of the compression algorithms used today (in videoconferencing and digital media such as DVDs) was feasible at that time. It was therefore assumed that a videoconferencing service would require something in the region of 1 Mbps to deliver acceptable video quality to users. The difference between the bandwidth required for a video call and a digitized-voice call was therefore large (about 12-to-1). There was no clear plan for how a switch would handle a mixture of traffic with such large differences in the bandwidth of video, voice, and data calls.

By 1971 the view had emerged that the standard unit of data bandwidth for data carried over the new network should be 64kbps, so that the same switching paths and transmission paths could be used for voice and data. At that time, 64kbps seemed like an excessive amount of bandwidth for any imaginable teletypewriter or screen-based data application.

In retrospect, 64kbps turned out to be a good choice. This is now the absolute minimum that Internet users expect in worst-case situations (dial-up internet connections or data connections via satellite telephone data). With improvements in video compression, and the development of methods for "inverse multiplexing" – combining several lower bit-rate channels into one higher bit-rate channel – by 1995, fairto-good quality videoconferencing was able to work over ISDN connections, using two, four, or six 64kbps connections, depending on the desired video quality.

The idea of a switched data network, that is, the data equivalent of a telephone network, in which any end-point device could connect to any other, based on an "address" (the generalization of the concept of a telephone number), was not new. The telex network, which grew up in the 1940s, was essentially a low-speed data network. Early work on data networks derived from this. But in the 1960s work on data networking, undertaken in several countries, started to branch into two separate lines of research:

- A "circuit switching" approach (like telex, or ISDN), in which a specific path through the network is selected and earmarked for the duration of the data "call". At the end of the call, the path (also called the "circuit" or "connection") is broken down.

- A "packet switching" approach (like the Internet), in which no attempt is made to establish a "call" or "connection" between the two end points. Instead, information is broken up into data packets which are each labeled with the target end point address and which make their own way through the network to their destination, sometimes being carried via different routes as traffic conditions in the network change, and sometimes arriving in a different order to that in which they were launched into the network. (As a result, packets need to be numbered and a mechanism at the receiving end has to sort them into the correct order before extracting the information from the packets.) The packet switching approach was thought to offer several advantages for data communications, such as efficient use of network capacity (bandwidth) and resilient operation in the face of failed network components.

During the late 1960s, at Dollis Hill and in various groups based in the City of London, the mainstream work in the BPO was on circuit switching. There was, however, a separate team at Dollis Hill working on packet switching – although at the time this was seen as long-term research that would not immediately lead to creation of a commercial network. The packet switching group was sharing its results with the US DARPA-funded work on the Arpanet (the forerunner of the Internet).

The most advanced countries in data networking in 1970 were Norway, Sweden, and Denmark, who had worked together to build a circuit-switched data network, separate from the telephone network. This was known as the Nordic Public Data Network (NPDN). It was the first circuit-switched data network to be operated commercially and it was used by a number of industries in these countries, including the banks. The BPO had been working closely with the Nordic group and had agreed to conduct a UK pilot network as an extension of the NPDN in the UK.

By 1970, work on switched data networks (having branched into circuit-switched and packet-switched lines of research and development) became further fragmented. In both the US and UK, work on packet switching started to branch into two distinct lines of development:

- ***Pure packet switching***. This was the continuing DARPA-funded work on packet switching which, twenty-five years later, would come to dominate data communication in the form of the Internet. This work aimed to give all the switching systems (or, as they later came to be called, "routers") in the network the ability to independently select the next leg of each individual packet's journey to its target address, based on information received from other routers about the state of the network and its own assessment of the state of its links to other routers.

- ***Blended circuit/packet switching***. This was a blended approach in which data is conveyed in packets, but the exact path to be taken by the packets between two end points was mapped out at the start of a data "call" and registered as a Switched Virtual Circuit (SVC). This took away some of the benefits of pure packet switching (namely, the automatic optimization of traffic levels on each link as a result of each individual packet being routed via the least-congested link at any moment in time, and the ability of the network to route packets around failed links, making the network very resilient). However, the SVC approach greatly reduced the processing load on each switching system in the network, because routing decisions need to be made only once per call instead of once per packet. The SVC approach was believed, by its designers, to be the only way of building an economically viable network using components that were available in the early 1970s.

In fact, the supporters of the SVC approach were right about the load that packet switching would place on a switching system. The amount of processing power and memory that would be needed to support packet switching in a commercial public data network seemed like science fiction when discussed in 1970. It was not until over twenty years later that vendors like Cisco would be able to buy the necessary chips to allow them to build affordable "routers". In 1970 nobody correctly predicted the staggering cost/performance improvements in memory and processor chips that would take place in the next twenty years as a result of the consumption of these chips by the computer industry, particularly the personal computer industry. In 1970, a rational electronics engineer would have predicted that powerful-enough chips, at low enough prices to support the commercial use of packet switching, would not be available until 2010 or later.

Development of a commercial SVC-type packet switching system was not attempted in the UK; but in the USA a few groups of determined engineers pressed ahead and created systems that were placed into service. The most well known company in the field was Telenet, which built switches and operated the Telenet public data network.

In 1971 the international sales team from Telenet started trying to interest telephone administrations in Europe in their products. The sales team lavishly entertained a number of senior members of the BPO and, rather abruptly by the slow-moving standards of the BPO, all work on the circuit switched data network was stopped. An announcement was made that the BPO would be deploying a public data network, to be known as the Packet Switched Service (PSS), using Telenet's packet switches.

This left the BPO teams who were working on what became System X feeling uncertain about whether the future network would be required to carry data as well as voice, or whether packet switching or circuit switching was the ultimate solution for data communication in the UK.

Partly as a result of the confusion caused by the PSS announcement, and partly because they had plenty of other things to think about, the Dollis Hill team gave very little thought to the requirements of carrying data on the new network for the next ten years.

Interest in an integrated voice/data network continued in other organizations in Europe, including ITT and Ericsson. The biggest technical problem that these organizations recognized was how to get the data service into homes and offices without requiring huge investments in new "last mile" infrastructures (between the local exchange and the home/office). Field trials of various approaches had shown that it might just about be possible to achieve digital transmission at around 120 kbps to 160 kbps over a normal copper-pair local line. This led to the idea that a standard could be developed that would support two 64kbps channels – one for voice and one for data – plus a lower speed channel (16 kbps) for signaling. At the local exchange the two 64kbps could be connected to the same digital switching system and this system would treat telephone calls and data calls in a similar way. The standard that eventually emerged (after about ten years of development and testing) was what is now called the ISDN basic rate interface, or "2B+D" interface (where B represents a 64kbps data channel and D represents the signaling channel).

Larger business users would not have to depend on the 2B+D interface to connect to the network with ISDN capability. Higher speed connections, starting with the confusingly named "primary rate interface" (easily confused with "basic rate interface"), provide many 64kbps channels over a dedicated digital circuit. However, by themselves these higher-speed interfaces would be of limited use to companies that wanted to exchange data with residential users and smaller businesses. It was therefore essential to provide ISDN services to smaller businesses and residential users, via the 2B+D interface, in order to bring to the network the traffic that would terminate on the higher-speed connections used by large companies.

Work on the 2B+D “last mile” solution, and work on the various protocols that would need to be defined for the user’s device (a terminal or computer) – to negotiate with the network the setting up of a data call – continued through the 1970s and into the early 1980s. The results of this work were presented at the ITU standards organization (at that time called the CCITT, now ITU-T). This led to the publication of the “I” series of recommendations, covering many aspects of ISDN, in 1984. The term “ISDN” emerged sometime around 1980-81. (Nobody has claimed credit for first use of the term, but it was probably one of the members of the various CCITT “I”-series committees.)

It took quite a long time to get from the publication of the standards in 1984 to a really dependable public service, even in the technologically advanced countries like the UK and USA. Early pilot services were started in the early 1990s, but it was not until about 2000 that interpretations of the CCITT recommendations reached a point where the service could be described as “plug and play”.

The slow pace of development of workable and widely adopted standards for the ISDN “last mile” gave British Telecom roughly ten years (starting in 1984) to retrofit System X with the necessary software and additional hardware modules required to support ISDN. So, while it is true to say that supporters of an all-digital design for System X used the idea of an integrated voice/data network to support their arguments throughout much of the 1970s (when they were losing ground to the supporters of the hybrid reed-relay/digital switch design), the actual impact of this argument on the final 1982 design of System X was minimal. And, as shown by the long time it took for reliable ISDN services to be deployed, System X in 1982 was only “ISDN ready” in the sense that it did not contain reed relays.

When a usable ISDN service did emerge in the late 1990s, many of the potential sources of traffic imagined in 1970 had found other solutions:

- First, dial-up connections over the telephone network had improved dramatically in speed and reliability, from 300bps to error-free 56kbps transmission, as a result of developments in modem technology and processor chips embedded in modems that could perform error correction and data compression. These modems were widely used by companies for traveling staff and for customers to access their computer-based services, such as electronic banking, database access, electronic mail, chat rooms, and so on. They were also used to gain access to the pre-Internet information service providers like Compuserve and AOL. (People tend to forget that there was a time when AOL was not even connected to the Internet.)

- Second, the packet switched services like PSS in the UK, Telenet and Tymnet in the USA, and others, had met the needs of users who wanted to communicate with several different computer systems in succession, without having to dial a new telephone call for each one. They also provided international data calls a lot more cheaply than using a modem over an international telephone call.

- Third, the Internet evolved from a network used by universities and research organizations into a global, commercial network which then became the preferred solution to almost all data communication problems from the late 1990s onwards. This evolution was driven by the emergence of:
    - The Domain Naming standard (Jon Postel, Paul Mockapetris, and Craig Partridge, 1986)
    - The concept of embedded links (hyperlinks) in displayed material (originally proposed by Ted Nelson in the mid-1960s)
    - The concept of screen pages built from multiple sources
    - The HTML and HTTP standards, which tied the above three things together in a practical solution to delivering combined text and graphics information to a screen-based display terminal (Tim Berners-Lee, 1989)
    - The early browsers that turned these concepts and standards into usable PC programs (starting with Mosaic, 1993)

> None of these five foundations of the Internet as we know it today was even dreamed of in 1970. Even in 1984, when the ISDN standards were finally published, ideas on what data communication could ultimately be used for were poorly developed. The need for standards (like HTML) that represented higher protocol layers than those discussed in the CCITT standards groups was not widely recognized.*

In spite of the alternatives to ISDN for data communication that arose while ISDN was in the process of being specified and implemented, ISDN played a limited but important role to play for a few years when it finally became a dependable service for high-speed dial-up access to the Internet, videoconferencing, providing a cost-effective backup path backup for leased circuits. However, these uses of ISDN are likely to cease in the near future.

The emergence of System X

Having covered, in the first half of this paper, the various new technologies and design concepts that emerged between 1949 and 1976, together with the technology decisions made by the BPO and the various pre-System X systems and pilot systems that were built and deployed, we will now go back to 1970 and look at the organizational and political events that led up to the formal start of System X development.

In 1970 Roy Harris had returned to work. Rather than being asked to resume leadership of the switching development team at Dollis Hill, he was asked to lead a new joint BPO/manufacturers organization that was be tasked with rebuilding the collaboration and trust that existed under JERC, prior to the dissolution of The Ring. The new organization was called the Advisory Group on Systems Definition (AGSD). Its headquarters were in Lambeth, near St Georges Circus.

Under Roy Harris's leadership, AGSD quickly evolved into a complex structure of committees, subcommittees, and subsubcommittees, debating every aspect of what was later to become System X. BPO engineers from Dollis Hill and the various development groups in the City of London all wanted to be on an AGSD team. They regarded an AGSD meeting as "a nice day out of the office". The pub close to AGSD became overwhelmed by increased lunchtime business from BPO and manufacturers. Conference rooms at the AGSD offices became very difficult to book on the days that the pub featured a lunchtime striptease act.

While the representatives of the manufacturers on the various AGSD committees were clearly technically competent, representatives from the BPO felt that the manufacturers were avoiding any real sharing of information and were trying to keep the meetings on the level of intellectual debate, rather than concrete design. The STC representatives were treated with the greatest coolness by the Plessey and GEC representatives because it was known that STC's parent, ITT, was working on its own SPC system for the global marketplace (System 12). The fear was that STC would take any research results that it picked up from GEC and Plessey at the AGSD meetings and pass them on to their colleagues in ITT Belgium.

To a first approximation, the hundreds of man-years of discussion and drafting of position papers that went on under Roy Harris's chairmanship of AGSD between 1970 and 1974 produced nothing significant

---

* In 1985 I addressed this topic in the *Telecommunications Users' Handbook* where I wrote: "Other examples of what is theoretically possible with interactive services include making travel reservations, booking theatre tickets, paying bills against a bank account or credit card account, seeing your up-to-date bank statement or credit card statement, and sending messages to other users (electronic mail). . . The use of personal computers for accessing [interactive services] will, I believe, become more or less the 'normal' method of access. . . Later on, the more demanding customers will want to see coloured charts, graphs, and maps on their screens. The simple information services will then add colour and graphics as options. For this to happen what is badly needed is a new [display standard] that gives 80 characters-per-line text and colour graphics with really fine detail – not primitive figures made up of blocks of colour. . . Such a standard might even include a means of sending the image of a photograph . . . built up out of 'dots' of colour like a newspaper photograph. . . At the time of writing it is not clear whether it will come from an international standards body, a major vendor like IBM, or from personal computer software producers." When I wrote this it would be another four years before Tim Berners-Lee defined HTML, and a further four years before the release of the first practical browser (Mosaic). I used the word "dots" because the term "pixels" (which is short for picture elements) was not in use in 1985.

that formed any part of the final System X design in 1981. This was partly a result of the mutual distrust and suspicion between the three manufacturers. But the principal reason that AGSD failed to create any useful basis for later System X development work was the lack of any clear goals for AGSD. (AGSD in fact had a formal charter, although this was extremely vague – it talked about defining modular subsystems in functional terms to a level of detail that would "provide a basis for subsequent development and manufacture" – but there was no agreement on what this meant in practical terms.)

Most of the AGSD work focused on how the overall System X design could be broken down into a number of subsystems, with clearly defined interfaces to one another. These subsystems all had acronyms, such as DSS, SIS, PU, PPU, and so on. Most of the representatives at AGSD were hardware engineers; and so the debates about these functional "black boxes" were very much from a hardware point of view. There was no discussion of how the software was going to be written, how many man-years of software would be required, how the programming teams would be organized, or how the control elements of each subsystem would be specified in a way that would be useful to a programming team.

The BPO representatives on AGSD had plenty of distractions outside of AGSD meetings. The whole of Dollis Hill was being relocated to a new campus at Martlesham Heath in Suffolk. R8 and R15 were among the first groups designated to move. The construction at Martlesham had fallen behind schedule, so R8 and R15 had to move into temporary offices at Eastgate House and Scottish Mutual House in Ipswich. The move of so many professional staff into a largely rural area had created an acute shortage of suitable housing, so staff who were relocating had to devote a great deal of time and energy to their personal moves. As a result of these distractions, very little new work was undertaken in R8 and R15 for most of 1971; and when work resumed, the efforts of the various sections seemed to have fragmented and lacked a clear common goal.

The only person who showed any grasp of the upcoming problem of writing the software for System X was John Jarvis, head of section R8.2. Jarvis was a hardware engineer but he asked a very penetrating question about software: "How complex are telephone switching control functions?"

Jarvis initiated a project called "The Quantification of the Relative Complexity of Control Functions". Jarvis did not create this project from any sudden insight into the problems of large-scale software development projects. Rather, it was a hardware engineer's reaction to a new problem that he sought to understand in engineering terms. In this project, staff were asked to guess how many man-weeks it would take to write the software to perform a particular well-defined function that was a small part of the overall function of controlling a telephone switching matrix. Two staff, who also had a hardware engineering background, but who had learned to program in assembler, then wrote the software. After demonstrating its correct operation on a small testbed system, they counted the lines of code they had written. (They also kept track of how long it had taken to write this code.) Finally, they compared the results with the original guesses.

The conclusions that Jarvis drew from this work were interesting:

- Even educated guesses about the complexity of functions are very often wildly wrong: what appears to be simple in concept turns out to be very complex to program, and vice versa;

- The complexity of the functions performed in strowger and crossbar exchanges by "a bunch of relays and selectors" turns out to be frighteningly high; and

- The total effort required to complete an SPC system in any reasonable timeframe greatly exceeded the programming resources in R8 and R15 at that time.

Unfortunately, Jarvis's work was ignored by BPO Management and derided by the self-appointed gurus of software in the other sections. Aside from this account that you are reading now, the work of Jarvis and his team has probably never been mentioned again.

Processor selection

In April 1972, BPO Management decided that it was important that a specific computer system be selected for System X, so that software development work could be directed at a particular physical computer architecture. In the previous year the view had emerged that no standard commercial computer system (such as the Honeywell 516) would be reliable enough on its own for use in a telephone exchange. Various schemes for using two commercial machines in a main/standby configuration had been considered and rejected. The consensus was that a multiprocessor architecture would be best, with many processor units and memory units sharing a common bus. Tasks would be distributed across the various processors – with some mechanism for detecting, and removing from the active pool of devices, any faulty processor unit or memory unit. Such systems were under development in various organizations around the world (although, in the end, very few of these developments resulted in systems that went into production).

Both Plessey and GEC had been doing work on multiprocessor systems and claimed to have designs that could form part of System X. It was decided by BPO Management that a very rapid evaluation of these two systems should be carried out, during April and early May 1972, and one of them selected.

A "Processor Selection Committee" of about thirty BPO staff was formed to study the two contenders – the Mark 2BL multiprocessor system from GEC, and the System 250 multiprocessor system from Plessey. The Plessey system was a "real" system, in that it had been built and deployed as the basis of the Ptarmigan battlefield communication system. The GEC system was a "paper" system whose component parts had been bench-tested, but which had not been built or tested as a complete working system.

The Committee spent about four weeks reviewing the two designs. They visited Plessey's Taplow facility (17-18 April, 1972) and GEC's main site at Coventry (27-28 April) to see the prototypes and speak to the development staff. The Committee wrote a report recommending the choice of the Plessey System 250, largely because they felt that it was a real, proven system – in contrast to GEC's paper design. However, for reasons that were never disclosed, the recommendations of the selection committee were overturned by BPO Senior Management and a choice of the GEC Mark 2BL was announced.

Regardless of the way that the choice of the 2BL came about, the selection process was in any case premature. It would not be until around 1977 that the creation of the System X software would reach the coding stage, requiring programmers to take into account the target machine on which their software would run. So, it might have been better had GEC and Plessey been asked to work towards creating prototype machines (by some date in the future) for use in a side-by-side comparison.

The formation of TSSD

After four years of AGSD committee meetings between 1970 and 1974, and fragmented research efforts in R8 and R15, BPO Management recognized that the effort to build a family of SPC telephone exchanges (to take over from the stopgap TXK and TXE systems) was stalled.

Roy Harris, feeling that AGSD did not represent the sort of power base that he needed to get things moving, had pitched to BPO Senior Management the idea of a dedicated team (headed by himself) that would steer the development of the new family of systems and would manage a set of outsourced development activities to be performed by the manufacturers. He asserted that the "black box" subsystems defined in the AGSD deliberations could form the basis of separate blocks of development work contracted out to different manufacturers. Reportedly he had in mind a core team of about 30 of the best engineers drawn from R8, R15, and various development groups in London.

Senior Management decided to go along with Harris's proposal. In April 1974 they announced the creation of a new organization, called TSSD (Telecommunications Systems Strategy Division), to be headed by Harris. However, instead of creating a small, select team as suggested by Harris, they made

TSSD into a small empire, putting the whole of the former R8 and R15 into TSSD, along with all the development groups in London that were working on any aspect of telephone switching and signaling. This immediately created an unwieldy organization that Harris never brought under control.

It was in this announcement that the name “System X” became the official name of the system that was to be built. At Dollis Hill in the late 1960s the system that R8 and R15 were planning to build had been called “ADMITS”, which stood for “Adaptable Dispersed Modular Integrated Telecommunications System”. In 1971 this name had been dropped and the work of R8 and R15 was renamed the “ETP”, or “Evolutionary Telecommunications Project”. The intended output of the ETP no longer had a name, but for the year or so prior to the formation of TSSD the name “System X” had come to be used sarcastically by BPO staff – as a comment on the fact that Senior Management were so unclear about the direction of the Research Department work that they had not even been able to agree on a name for it. In one of the few clever decisions surrounding the project, Senior Management decided to defuse the joke by adopting “System X” as an official system designation. (Later, the specific versions of System X for different functions within the network were labeled with the prefix TXD, which followed the same theme as TXK and TXE.)

Harris had hoped that enough work had been done in “defining the subsystems” under AGSD to enable him to use the AGSD work as the specifications for development contracts, placed with the three manufacturers. However, on examination, it was found that the AGSD documents were too conceptual for this purpose. It was therefore decided to split the development work into two phases. The first would be a series of “Definition Contracts”, and these would be followed by a series of “Development Contacts”. In essence, the Definition Contracts would create functional specifications for System X. Or, to put it another way, *the BPO would pay the manufacturers to tell the BPO what the BPO wanted them to build*.

As an approach to developing a system so critically dependent for its success on the software components of the system, this was a recipe for disaster. What made it even worse was that TSSD staff would eventually find that they had minimal control over the progress and outcome of the Definition work. This led to the output of the Definition work falling far short of what could be considered a useful functional specification. (In fact, the creation of documents that more closely resembled functional specifications did not start until about two years later.)

The processor selection episode in 1972 had raised the level of mutual distrust and suspicion between GEC and Plessey, and things continued to deteriorate in the following two years. As a result, once the structure and general approach for System X development was announced, it proved difficult to get the process started. The period from June 1974 to July 1976 was taken up with legal wrangling over the wording of the Definition Contracts. Thus, the definition work that was supposed to have started in July 1974 did not actually start until July 1976.

Work starts under the Definition Contracts

After the two years of legal wrangling, work finally started on the first “Definition” phase of creating System X in July 1976.

The way that the various subsystems of System X would be divided between the three manufacturers had been locked down at this time. The processor subsystem, which came in two sizes – the PU (Processor Utility) and the smaller PPS (Pre-Processor Utility) – went to GEC. The Digital Switching Subsystem (DSS) and the Signaling Interworking Subsystem (SIS) went to Plessey. The Subscriber Switching Subsystem (SSS) – the reed-relay switching matrix for the local exchange – went to GEC. The Message Transmission Subsystem (MTS) and the Maintenance Control Subsystem (MCS) went to STC.

The allocation of the definition, design, and implementation of these subsystems between the three manufacturers was (like the splitting of the complete system into subsystems in the first place) based on a predominantly hardware view of the project. If the project had been, for example, a transmission system with no software component, such a division of work would have been tolerable (though not ideal, since it would make communication about the detailed interfaces between the subsystems more difficult than if the work were to take place within a single company). But because System X depended critically on its software components, the way that the subsystems were allocated between the manufacturers was a disaster. It meant that, for example, Plessey staff were designing SIS software that would control hardware components being designed by GEC.

Later on, once coding started, it also meant that Plessey and STC would be writing software to run on a processor system designed by GEC. This, by itself, would not have been a problem if the design of the 2BL had been completed and a near-production version of the processor system had been available for software testing. However, the 2BL design work continued as GEC tried to move it from a paper design to a real design. Plessey and STC would later complain bitterly that they could not get the information, about the detailed operation of the 2BL, needed to make any progress on coding and testing.

Another feature of the way that the subsystems were assigned to the manufacturers was the poor alignment between the assigned subsystems and each manufacturer's existing areas of expertise. This was reportedly because of lobbying by the manufacturers to get the subsystems in which they were least expert so that, in working on these subsystems, they would round out their own expertise. It is unclear why the BPO went along with such a plan that was clearly not going to make for a speedy and efficient development process. This is yet another illustration of how the manufacturers cynically disregarded the needs of their customer (the BPO) and regarded the money that the BPO paid to them as a Government hand-out to be used for their own benefit.

What made the way that the subsystems was divided up still more complicated was the fact that the System X work was "multi-site, with a vengeance". The main core of TSSD was in Ipswich. Since the move of R8 and R15 to Ipswich in 1971, a third temporary building was rented to accommodate the growing number of staff. Other parts of TSSD were in Central London, in what had previously been Development groups. Plessey's System X staff were split between Poole, Liverpool, and Taplow. STC's System X staff were in New Southgate. And GEC's System X staff were in Coventry. The System X development organization was thus split over a total of nine buildings, most of them in different cities.

The way that the definition work was supposed to proceed was that teams (or, in a few cases, individuals) within TSSD were each given responsibility for "supervising" one definition contract. It was their job to meet with the manufacturer's team periodically, review drafts of definition documents and other deliverables, provide feedback to the manufacturer, request that the manufacturer make changes to their documents in terms of style, content, or general approach, and report progress to Senior Management. However, within just a few months (in Autumn 1976) it had already become clear to the TSSD teams that their control over the direction of the manufacturers' work was minimal. Teams that had reported to Senior Management that they were unhappy with the direction taken by the manufacturers, or were unhappy with the way in which the definitions were being written, quickly found that Senior Management were not prepared to assume any reasonable measure of control over the work of the manufacturers. (I tried, and failed, to reject a major deliverable from one of the manufacturers and require it to be substantially revised. My manager told me that I could not do this because such an action would have "political repercussions". He ordered me to accept it and not "rock the boat".)

Clearly there was a fear that the manufacturers and their lawyers, having delayed the start of the definition work by two years, might start to become difficult again and cause progress to grind to a halt. Thus, the tone was set for members of TSSD that, far from being in the driving seat on System X, they were just there to "rubber stamp" the monthly work acceptance sheets so that the manufacturers could receive their considerable fees for the definition work.

With this as the starting point of the next five years' work, it is not surprising that progress was as slow as it turned out to be. Worse still, all the difficulties created by the management structure for the project, and the subdivision of the work between groups in many different locations, were exacerbated by the fact that none of the participants in the project, at the management level or the programmer level, had any experience of writing a large, complex application of any sort, let alone a real-time control system. These factors, taken together, created one of the least effective SPC system development projects in the world.

The total man-years consumed on System X software development between 1976 and 1981 was more than an order of magnitude greater than the equivalent efforts in AT&T, Ericsson, Philips, or any of the other major manufacturers around the world that completed the development of an SPC system.

As is common in software development projects that start to get behind schedule, the manufacturers soon started to add more and more programmers. This created the usual diminishing returns to scale, and eventually negative returns. In 1979, in a press release about System X, the BPO said that one of the reasons that System X development was so far behind schedule was “the shortage of programmers”. It said that programmers from software houses had been drafted in because the manufacturers had not been able to fill all their open positions for programmers.

It is hard to accurately gauge the level of effort that was expended over the period 1976 to 1981 to develop the software for System X because, in the panic that ensued once it was recognized that the software would not be ready in 1978, there were so many different organizations contributing staff to the effort that an accurate count of man-years was never attempted. However, it seems likely that the true figure would lie somewhere in the range 3,000 to 6,000 man-years, including all phases of software development (requirements definition, design, coding, and testing).

If we look at comparable efforts in other countries, it is clear that this level of effort was outrageously high for the task at hand, by a factor that far surpasses any of the horror stories told about large software development projects in Frederick Brooks’s book, *The Mythical Man-Month*, published in 1974 (a book that BPO Senior Management would have done well to read). The work on SPC systems at AT&T’s Bell Labs between 1963 and 1976 had shown that the total effort to create a full suite of software for an SPC system, covering local and trunk exchange functions, could be as high as 300 man-years (counting false starts, re-works, and experiments in how to organize the programming team). The team at Ellemtel (an offshoot of Ericsson jointly owned with the Swedish telephone network operator) who developed the AXE10 were better organized. They had the benefit of experience with a first-generation Ericsson SPC system (the AKE series). They also hired staff from other countries, including the USA, and they took the time to read papers that had been published on SPC work in the USA and Japan. Based at a single site on the outskirts of Stockholm, they were able to create the software for the AXE10 with an effort of less than 100 man-years.

Clearly the lessons of “The Mythical Man-Month” were re-learned at huge expense to the UK telephone subscriber between 1976 and 1981.

System X is completed

An early System X pilot system went into service on 1 July, 1980 at Baynard House exchange in London. This system performed the simple function of a tandem exchange (an intra-metropolitan-area trunk exchange) – exactly the same function as the Empress pilot in 1968. Although the switching matrix of this System X exchange was built using fairly new integrated circuits, functionally it differed very little from that of Empress. Thus, the only major difference between Empress tandem and the Baynard House System X tandem was that the hard-wired electronic logic of Empress had now, *after twelve years*, been rendered as software running on a processor system. The Baynard House cut-over was probably the most vivid illustration of how the UK’s *eight-year lead* in digital switching technology had been turned into a *four-year lag* by disorganized and inexperienced attempts at managing the software development component of System X.

The main reason for choosing a tandem exchange for this pilot was that the more complex parts of the software needed to control a local exchange were not ready. The first System X local exchange to go into service was at Woodbridge, Suffolk (near the Martlesham Heath campus) in 1981.

Following the Baynard House and Woodbridge pilots, the widespread deployment of System X proceeded slowly at first. Fewer than 500 System X exchanges were installed over the next five years.

In November 1986, the first Ericsson AXE10 (named, for the purposes of UK deployment, "System Y") was placed into service at Sevenoaks. From this point on the installation of System X and System Y exchanges proceeded in parallel. By 1995 there were 2,000 System Y exchanges and the number of System X exchanges had increased by 4,000 to a total of 4,500. The "X+Y" deployment was not completed until 11 March, 1998, when the last two TXE4s were replaced, respectively, by a System X at Leigh-on-Sea and a System Y at Selby. At this point the UK network became "totally digital", fulfilling Tommy Flowers's vision in the 1950s of an all-electronic telephone network.

Not only was System X not used for 100% of the exchanges in the UK, as originally intended, but attempts to make System X into an international product failed miserably. A joint BT/manufacturers marketing organization was set up in 1979, led by John Sharpley, but this was later wound up when it became clear that all major countries that did not have their own telecommunications equipment industry had already selected an SPC system from products, such as the Ericsson AXE10, that had been available for several years. (In this respect, the development of System X parallels that of Concorde, which also failed to find the buyers that were originally expected in countries other than the UK and France.)

Summary

The development of System X was not a failure, in as much as it was eventually completed. It was also not the worst-managed SPC system development in the world – an award that would go to ITT for System 12, a development project so out of control that its horrendous development costs were a major factor in the total financial failure of the ITT empire. The reasons for System 12's timetable and budget overruns were similar to the reasons for System X's. Like System X, System 12 suffered from a complete failure to properly manage the software development process. System 12 was also excessively ambitious in trying to convert networks to being fully digital. The complete System 12 "solution" included digital telephone instruments and ISDN-type digital local lines between the subscriber's premises and the local exchange.

But although System X was not a complete failure, its development was hideously expensive – over a billion pounds. This cost, which got the UK industry to the point where the ordering of approximately 4,600 System X exchanges could start, could have been avoided. The fact that the Ericsson AXE10 was installed in parallel with System X from 1986 onwards clearly demonstrates that the BPO did have an alternative to proceeding with the development of System X in 1976, since the Ericsson AXE10 was close to going into full production that year. Discussions with Ericsson in 1976 about licensing the AXE10 design to the UK manufacturers would have been a viable (although politically contentious) alternative to proceeding with System X.

The tremendous cost and timetable overruns of System X resulted from gross errors of judgment by senior technical managers in the BPO. In particular:

- They failed to recognize that the "control" function of a telephone exchange was very complex. They reasoned, incorrectly, that the control functions in a telephone exchange must be fairly simple because they are performed in electromechanical exchanges (strowger and crossbar) by simple relays and selectors. They failed to see that a large number of simple relays and selectors, richly interacting with one another, represented a great deal of "intelligence" that would require a huge software development effort to reproduce in an SPC exchange.

- They failed to talk to anyone who had worked on SPC systems in AT&T and thus missed a valuable opportunity to gauge the size of the effort required to create software to control a telephone exchange, or to learn about the challenges of organizing a team of programmers of the size needed for the task of writing System X's software.

- They failed to bring into the BPO, either as a fulltime member of staff or as a consultant, anyone with hands-on experience of managing a major software development effort.

- They ignored the few members of TSSD who tried to argue in favor of a process to gauge, at the start, the size of the required software development effort and to introduce a structured approach to the software development process.

- They assumed, without any analysis, and in most cases incorrectly, that the obvious subdivisions of the hardware design of the system (based on work at AGSD) would also make good subdivisions for the software components that worked with those subsystems.

- They outsourced the “definition” work (that is, the creation of the functional specifications for System X) to the manufacturers who were going to develop the system so that, in effect, the manufacturers had to guess what it was the BPO wanted, based on the very general descriptions created by AGSD.

- They did not define a coherent overall architecture for the system, nor did they empower anyone in the BPO team to make design decisions that overarched the work of the individual teams within the manufacturers.

- They allowed the software development work to proceed in the resulting fractured manner across many geographically dispersed sites.

- They failed to take an appropriate leadership role in directing the work of the dispersed teams.

All this is now history. But it provides valuable lessons in how *not* to organize a project with a major software component. It underscores the importance of a structured approach, strong project leadership, quantification of the size of the task at the start, and a clear vision, from project initiation, of how the overall task is going to be accomplished.