[From: Chapter 25 in *The Social Shaping of Technology*, ed. D. MacKenzie and J. Wajcman (Buckingham, UK: Open University Press, 1999), 2nd edition)]

Cold war and white heat: the origins and meanings of packet switching

Janet Abbate

Of all the technical innovations featured in the ARPANET, forerunner of the Internet, perhaps the most celebrated was packet switching. Packet switching was an experimental, even controversial method for transmitting data across a network. Its proponents claimed that it would increase the efficiency, reliability and speed of data communications, but it was quite complex to implement, and some experts argued that the technique would never work. Indeed, one reason the ARPANET became the focus of so much attention within the computer science community was that it was the first large-scale proof of the feasibility of packet switching. The successful use of packet switching in the ARPANET and other early networks paved the way for the technique's widespread adoption, and at the end of the twentieth century packet switching dominates networking practice.

To many computer professionals, packet switching appears to have obvious technical advantages over alternative methods for transmitting data, and they have tended to treat its widespread adoption as a natural result of these advantages. In fact, however, the success of packet switching was not automatic: it had to be socially constructed. For many years there was no consensus on what packet switching actually was - what its defining characteristics were, what advantages it offered, how it should be implemented. Before the technique could achieve legitimacy in the eyes of data communications practitioners, its proponents had to prove that it would work by marshalling the resources to build demonstration packet switching networks. The wide disparity in the outcomes of these early packet switching experiments demonstrates that the concept could be realized in very different ways, and that, far from being a straightforward matter of superior technology winning out, the 'success' of packet switching depended greatly on how it was interpreted.

Packet switching was invented independently by two computer researchers working in very different contexts: Paul Baran at the Rand Corporation in America and Donald Davies at the National Physical Laboratory in England. Baran was the first to explore the idea, around 1960; Davies came up with his own version of packet switching a few years later and subsequently learned of Baran's prior work. Davies was instrumental in passing on the knowledge of packet switching that he and Baran had developed to Lawrence Roberts, who oversaw the creation of the ARPANET. However while Baran's and Davies's versions of packet switching had some basic technical similarities, their conceptions of what defined packet switching and what it was good for were very different. Much of this difference was due to the strong political pressures that were brought to bear on computing research in Britain and the United States. Large computer projects in both countries were developed in a context of government funding and control, and national leaders saw computers as a strategic technology that was vital to pursuing important political goals. But in the very different policy contexts of the United States and United Kingdom, packet switching took on different meanings for Baran, Davies and Roberts. Packet switching was never adopted on the basis of purely technical criteria, but always because it fitted into a broader socio-technical understanding of how data networks could and should be used.

NETWORKING DR STRANGELOVE: THE COLD WAR ROOTS OF PACKET SWITCHING IN THE UNITED STATES

In 1964, movie theatres across the United States presented Stanley Kubrick's brilliant black comedy of cold war paranoia, *Dr Strangelove.* The film, though humorous, highlighted a serious problem for American defence strategists: the vulnerability of communications channels to disruption by a Soviet attack, which might make them unavailable just when they were needed most. In the movie, a psychotic air force commander named Jack D. Ripper sets in motion a nuclear holocaust by invoking a strategy of mutually assured destruction called 'Plan R.' Plan R - which allows Ripper to circumvent the president's authority to declare war - is specifically designed to compensate for a wartime failure in command, control and communications. As the movie dramatizes, this plan is hardly ideal, since it allows Ripper to launch a 'retaliatory' attack even though no Soviet first strike has actually occurred. In reality (as the film's disclaimer states), the air force never had any such strategy. In fact, the air force was at this time exploring a very different solution to the problem: building a communications system that would be able to survive an enemy attack, and thus maintain 'proper command and control.'

The need for 'survivable communications' was a generally recognized problem by the early 1960s, and among those intent on solving it was a researcher at the US air force's premier think-tank, the Rand Corporation. Founded by the air force in 1946 as an outgrowth of operations research efforts initiated in World War II, Rand (or RAND, an acronym for Research and Development) was a non-profit corporation dedicated to research on military strategy and technology. Rand was primarily funded by contracts from the air force, though it served other government agencies as well. The corporation attracted talented minds through a combination of high salaries, relative autonomy for researchers, and the chance to contribute to policy decisions of the highest importance.

In 1959 a young engineer named Paul Baran joined Rand's computer science department. Immersed in a corporate culture focused single-mindedly on the cold war, Baran soon developed an interest in survivable communications, which he felt would decrease the temptation of military leaders to launch a pre-emptive first strike:

> Both the US and USSR were building hair trigger nuclear ballistic missile systems ... If the strategic weapons command and control systems could be more survivable, then the country's retaliatory capability could better allow it to withstand an attack and still function; a more stable position. But, this was not a wholly feasible concept because long distance communications networks at that time were extremely vulnerable and not able to survive attack. That was the issue. Here a most dangerous situation was created by the lack of a survivable communication system.

Baran was able to explore this idea without an explicit contract from the air force, since Rand had a considerable amount of open-ended funding that researchers could use to pursue projects they deemed relevant to the American defense concerns.

Baran began in 1959 with a plan for a minimal communications system that could transmit a simple 'Go/No go' message from the president to the commanders using AM radio broadcasts. When Baran presented this idea to military officers, however, they immediately insisted that they needed greater communications capacity. So he went back to the drawing board and spent 1960-2 formulating ideas for a new system that would combine survivability with high communications capacity. Baran envisioned a system that would allow military personnel to carry on ordinary voice conversations or use teletype, facsimile, or low-speed computer terminals under wartime conditions.

The key to this new system was a technique Baran called 'distributed communications'. In a conventional communications system such as the telephone network, switching is concentrated and hierarchical. Calls go first to a local office, then to a regional or national switching office if they need to connect beyond the local area. Each user is connected to only one local office, and each office serves a large number of users. This means that destroying a single local office would cut off many users from the network. By contrast, a distributed system would have many switching nodes and many links attached to each node. Such redundancy would make it harder to cut off service to network users.

Baran described this new type of network in a series of eleven reports called *On Distributed Communications.* His proposed system (a small segment of which is shown in Figure 25.1) consisted of a set of several hundred switching nodes, each connected to other nodes by up to eight lines. There was also a set of several hundred multiplexing stations that provided the interface between the users and the network. Each multiplexing station would be connected to two or three switching nodes and up to 1024 users with data terminals or digital telephones. The switching was distributed among all the nodes in the network, so that the enemy could not disable the whole network by targeting a few important centres. To make the system even more secure, Baran planned to locate the nodes in areas remote from the network users, since population centres were considered military targets, and he designed the multiplexing stations with a wide margin of excess capacity, on the assumption that enemy attacks would cause some equipment to fail. He also added military features such as cryptography and a priority/precedence system that would allow high-level users to pre-empt messages from lower-level users.

To move the data through the network, Baran adapted a technique known as 'message switching' or 'store and forward' switching. A common example of message switching is the postal system. In a message switching system, each message (such as a letter) is labeled with its origin and destination addresses, and it is then passed from node to node through the network. The messages are temporarily stored at each node (such as a post office) until they can be forwarded to the next node or the final destination. Each successive node uses the address information to determine the next step of the route. In the 1930s, message switching began to be used in telegraph systems, with messages being stored on paper tape at each intermediate station before being transmitted to the next station. At first the messages were switched manually by the telegraph operators, but in the 1960s telegraph offices began to use computers to store and route the messages.

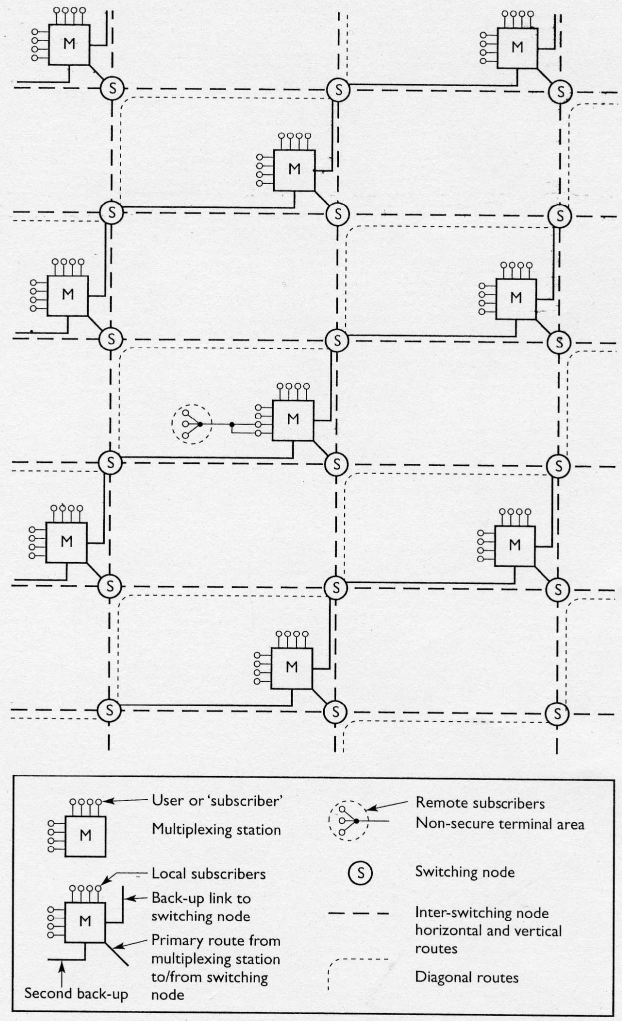

Figure 25.1 Baran's distributed network of switching nodes and multiplexing stations (Paul Baran 1964a *On Distributed Communications*, vol. VIII, Figure 1)

The postal and telegraph systems adopted message switching because it was more efficient than transmitting messages or letters directly from source to destination. Letters are stored temporarily at a post office so that a large number can be gathered for each delivery route. In the case of the telegraph, message switching also addressed the problem of uneven traffic flow on the expensive long-distance telegraph lines. If traffic was light, the lines would be underused and the excess capacity wasted; on the other hand, if the lines were overloaded, there would be a risk that some of the messages would be lost. Storing the messages at intermediate stations made it possible to even-out the traffic flow: if a line was busy, messages could be stored at the switch until the line was free. In this way message switching increased the efficiency, and hence the economy, of long-distance telegraph transmission.

Baran appreciated the efficiency offered by message switching, but he also saw the technique as a way to make his system more survivable. Since the nodes act independently in processing the messages and there are no pre-set routes between nodes, the nodes can adapt to changing conditions by picking the best route at any given moment….This increases the ability of the system to survive an attack, since the nodes can re-route messages around non-functioning parts of the network. Baran realized that survivability depended on more than just having redundant links; the nodes must be able to make use of those extra links: 'Survivability is a function of switching flexibility'. Therefore, his proposed network would be characterized by distributed routeing as well as distributed links.

Departures from other contemporary systems

Baran was not the first to propose either message switching or survivable communications systems to the military. Both types of system already existed or were in development, and there was one large-scale distributed communications network being constructed in the early 1960s. This was the AUTOVON network, designed and operated for the Defense Department by the US telephone giant AT&T. In 1961 AT&T had provided the army with a communications network called SCAN (Switched Circuit Automatic Network), and in 1963 they provided a similar network for the air force called NORAD/ADS (North American Air Defense Command/Automatic Dial Switching). The Defense Communications Agency, which was charged with coordinating the provision of communications services throughout the armed services, decided to integrate these networks into a new system called CONUS AUTOVON (Continental United States Automatic Voice Network). AUTOVON was not a message switching system but rather a special military voice network that was built on top of the existing civilian telephone network. AUTOVON had ten switching nodes and came on line in April 1964.

While AUTOVON was an example of distributed communications, Baran's approach differed from AT&T's in two crucial ways. First, although AUTOVON had nodes distributed throughout the system, control of those nodes was concentrated in a single operations centre, where operators monitored warning lights, analysed traffic levels, and controlled system operations. If traffic needed to be re-routed, it was done manually: operators at the control centre would make the decision and then contact the operators at the switching nodes with instructions to change routes. In Baran's network, control was fully distributed, as noted above. Nodes would be individually responsible for determining routes, and would do so automatically, without human intervention: 'The intelligence required to switch signals to surviving links is at the link nodes and *not* at one or a few centralized switching centers.' Clearly such a system would be more survivable than one dependent on a single operations centre, which, Baran noted, 'forms a single, very attractive target in the thermonuclear era'.

…The requirements of Baran's system would push switching and transmission technology to their limits, so it is understandable that contemporary experts reacted sceptically to his claims. The engineers in AT&T's Long Lines division, which ran the long distance telephone service and the AUTOVON system, tended to be familiar only with analogue technology, and they were sceptical of Baran's claims that an all-digital system could transcend the well-known limitations on the number of links per call.

Baran's system departed from traditional telephone company practice in other ways that show the effect of cold war military considerations on his design assumptions. For instance, the phone company tried to increase the reliability of the system as a whole by making each component as reliable as possible, and for an additional fee would provide computer users with lines that were specially conditioned to have low error rates. Baran chose instead to make do with lower-quality communications links and to provide redundant components to compensate for failures. Conditioned lines would be too expensive for a system with so many links, and in any case, the reliability of individual components could not be counted on in wartime conditions. As Baran observed, 'Reliability and raw error rates are secondary. The network must be built with the expectation of heavy damage anyway'.

Packet switching in Baran's system

Baran's proposed network began as a distributed message switching system. His final innovation was to alter message switching to create a new technique: packet switching. A message in his system could be anything from digitized speech to computer data, but the fact that these messages were all sent in digital form - as a series of binary numbers or 'bits' - meant that the information could be manipulated in new ways. Baran proposed that rather than sending messages of varying sizes across the network, messages should be divided into fixed-size units that he called 'message blocks'. The multi-plexing stations that connected users to the network would be responsible for dividing outgoing messages into uniform blocks of 1024 bits. A short message could be sent as a single block, while longer messages would require multiple message blocks. The multiplexer would add to each block a header that specified the addresses of the sending and receiving parties as well as other control information. The switching nodes would use the header information to determine what route each block should take to its destination; since each block was routed independently, the different blocks that made up a single message might be sent on different routes. When the blocks reached their destination, the local multiplexer would strip the header information from each block and reassemble the blocks to form the complete message. Baran's message block idea would eventually be widely adopted for computer networks. In these later systems the message blocks would be called 'packets', and the technique would become known as packet switching.

For all its eventual significance, the decision to transmit data as packets was not the original focus of Baran's work. As the title of his eleven volume system description, *On Distributed Communications,* indicates, he began with the idea of building a distributed network, an idea that had already been identified with survivability by people working in military communications. In describing the system, Baran tended to stress the idea of link redundancy, rather than other elements such as packet switching. But as he developed the details of the system, the use of message blocks emerged as a fundamental element. By the time he wrote the final volume of the series, Baran had changed the name he used to refer to the system to reflect the new emphasis: 'While preparing the draft of this concluding number, it became evident that a distinct and specific system was being described, which we have now chosen to call the "Distributed Adaptive Message Block Network", in order to distinguish it from the growing set of other distributed networks and systems'. What, then, was so important about packet switching? What did it mean to Baran and his sponsors?

Note that transmitting packets, rather than complete messages, imposed certain costs on the system. The interface computers had to perform the work of dividing users' outgoing messages into packets and reassembling incoming packets into messages. There was also the overhead of having to include address and control information with each packet (rather than once per message), which increased the amount of data that had to be transmitted over the network. And since packets from a single message could take different routes to their destination, they might arrive out of sequence, which meant that there had to be provisions for reassembling them in the proper order. All of this made the system more complex and presented more opportunities for failure. For Baran, these costs were outweighed by his belief that packet switching would solve some pressing problems and support some fundamental goals of the system.

Packet switching offered a variety of benefits. Baran was determined to use small, inexpensive computers for his system, rather than the huge ones he had seen in other message switching systems, and he was aware that, given the state of the art at the time, the switching computers would have to be simple in order to be both fast and inexpensive. Using fixed-size packets simplified the design of the switching node. Another advantage for the military was that breaking messages into packets and sending them along different routes to their destination would make it harder for enemy spies to eavesdrop on a conversation. But the biggest potential reward was efficient and flexible transmission of data. 'Most importantly', wrote Baran, 'standardized data blocks permit many simultaneous users, each with widely different bandwidth requirements^] to economically share a broad-band network made up of varied data rate links'. In other words, packet switching allowed a more efficient form of multiplexing, the sharing of a single communications channel between many users.

…In sum, packet switching was important to Baran because it furthered some key requirements of a survivable military system. Cheaper nodes and links made it economically feasible to build a highly redundant (and thus robust) network. Efficient transmission made it possible for commanders to have the higher communications capacity they wanted. Dividing messages into packets also increased security by making it harder to intercept intelligible messages. Packet switching, as Baran understood it, made perfect sense in the cold war context of his proposed system.

The US Air Force agreed with Baran's assessment and was eager to build a network based on his design. Internal Defense Department politics thwarted this plan, but Baran's ideas were widely disseminated among American military and civilian research institutes.

FORGING PACKET SWITCHING IN THE WHITE HEAT: NETWORKS AND NATIONALISM IN THE UNITED KINGDOM

While the United States was caught up in the cold war in the early 1960s, the United Kingdom was experiencing political upheaval of a different type. Just as the Americans were worried about a 'science gap' between their country and the Soviet Union, so there were widespread fears in the United Kingdom of a 'technology gap' with the United States. In 1963 Harold Wilson was elected leader of the Labour Party, at a time when that party, and much of the general population, felt that the country was facing an economic crisis. Politicians on all sides warned that the nation was falling behind the other industrial powers in its exploitation of new technologies, that there was a 'brain drain' of British scientists to other countries, and that the country's technological backwardness was at least partly responsible for its economic malaise.

Wilson addressed the technology issue head on in his address to the Labour Party's annual conference at Scarborough on 2 October 1963. Calling on labour and management to join in revitalizing British industry, Wilson stressed the importance of keeping up with the ongoing scientific and technological revolution, and he invoked a stirring vision of a new Britain 'forged in the white heat of this revolution'. When Labour came to power in the 1964 general election, Wilson was eager to act on his vision by implementing a new economic and technological regime for Britain. To oversee national technological development, Wilson created the Ministry of Technology, a major new department that assumed control of the Atomic Energy Authority, the Ministry of Aviation, the National Research Development Corporation (NRDC), and a number of national laboratories.

Mintech, as it came to be called, had two main aims: to transfer the results of scientific research to industrial development, and to intervene in industry so as to make private enterprise more efficient and competitive. Mintech was to have, in Wilson's words, 'a very direct responsibility for increasing productivity and efficiency, particularly within those industries in urgent need of restructuring or modernisation'. These industries included machines tools, aviation, electronics, shipbuilding, and - above all - computing.

One of the British scientists who took the lead in computing research was Donald W. Davies of the National Physical Laboratory (NPL) in Teddington, a suburb of London. The National Physical Laboratory had been established in 1899 to determine values for physical constants, standardize instruments for physical measurements, and perform similar activities involving standards and materials testing. NPL first became involved in computing in 1946, when a team at the laboratory, following a proposal by Alan Turing, built an early stored-program digital computer called the Pilot ACE. Davies had joined NPL in 1947 and had worked on the Pilot ACE; in 1960 he became superintendent of the division in charge of computing science, and in 1965 he was also named technical manager of the advanced computer techniques project. Davies was thus in touch with both the latest advances in computing technology and the government's plans to use that technology to aid the British economy.

If the watchword for Baran was survivability, the priority for Davies was interactive computing. He was particularly interested in a technique called time sharing, which allows many people to interact with a computer at the same time. During discussions with British and American colleagues, Davies became aware of a widely perceived obstacle to interactive computing: inadequate data communications. In early time-sharing systems, the terminals had been directly connected to the computer and were located in an adjacent terminal room. As time went on, people began locating terminals at some distance from the computer itself, either for the user's own convenience or, in the case of commercial time-sharing services, to offer access to customers over a wide geographic area. Distant terminals could be connected to the computer using dial-up telephone links and modems, but long-distance telephone connections were very expensive, and for data transmission they were also inefficient. Computer messages, as noted earlier, tend to come in bursts with long pauses in between, so computer users paid dearly for telephone connections that were idle much of the time. The high cost of communications put pressure on users to work quickly, sacrificing the user-friendly quality for which time sharing had been invented in the first place.

Davies was perhaps even more aware of the cost issue than his American colleagues. In Britain, unlike in America, there was no flat rate for local telephone calls. Also, while American researchers tended to think in terms of academic computing, where users normally accessed the machine from a relatively short distance, Britain had a larger percentage of users whose computer access came from distant commercial systems. Davies had a long-standing interest in switching techniques, and as he thought about the data communications problem, it occurred to him that a new approach to switching might offer a solution. He knew that message switching was used in the telegraph system to make efficient use of lines, and he believed that by adapting this technique to computer communications he could achieve similar economies. Like Baran, Davies came from a background in computing, rather than communications, so he felt free to suggest a technique that departed from traditional communications techniques but took advantage of advances in computer technology. Davies proposed dividing messages into standard-sized 'packets', and having a network of computerized switching nodes that would use information carried in packet headers to route the packets to and from time-sharing computers. He called this technique 'packet switching'.

Packet switching in Davies's system

Like Baran, Davies saw that packet-switching would allow many users to share a communications link efficiently. But he wanted that efficiency for a different purpose. Packet switching, in his view, would be the communications equivalent of time sharing: it would maximize access to a scarce resource in order to provide affordable interactive computing.

Davies presented his network ideas publicly for the first time at a talk in March 1966, which was enthusiastically received by an audience of people active in computing, telecommunications, and the military. It was one of the latter, an attendee from the British Ministry of Defence, who gave Davies the surprising news that packet switching had already been invented a few years earlier by an American, none other than Paul Baran. The fact that the military man knew about this earlier development when Davies himself did not underscores the very different contexts in which packet switching evolved in the two countries. Baran's foremost concern had been survivability, which was underlined by his use of terms like 'raid', 'salvos', 'target', 'attack level' and 'probability of kill' in describing the hostile conditions under which his system was expected to operate. Davies, on the other hand, did not *view* packet switching as a way to make the network survivable; after reading Baran he commented that 'the highly-connected networks there considered are not needed in a civil environment'. Instead, he thought the pressing need was for a network that could serve the users of commercial time-sharing services. The Labour government specifically wanted to redirect research and development efforts away from military projects and toward civilian industry. Davies shared this concern, which is evident in his plan to survey businesses to find out their data communication requirements. It also shows up in Davies's efforts to make the system easy to use. In his proposal for a national network he stressed that 'A further aim requirement we must keep in mind constantly is to make the use of the system simple for simple jobs. Even though there is a communication system and a computer operating system the user must be able to ignore the complexities'.

The main benefit of packet switching for this type of system was that it would bring down the cost of data communications. Davies found further meanings in packet switching, however, that derived from-his vision of a commercial network service. For instance, one of the merits Davies saw in packet switching was that it helped achieve fairness in access to the network. In an ordinary message switching system, each message had to be sent in its entirety before the next message could begin. In a packet switching system, users would take turns transmitting portions of their messages using time division multiplexing. If a user had a short message, such as a single command for a time-sharing system, the whole message could be sent in the first packet, while longer messages would take several time slots to transmit. The user with the short message would not have to wait behind users with long messages. This kind of fairness was appropriate for a system where computers were serving ordinary people's everyday needs, rather than transmitting life-or-death messages through a command hierarchy.

Davies also believed that packet switching technology could itself become a commercial product, and thus contribute directly to Wilson's plan to revitalize the British economy. In a 1965 proposal to have the General Post Office (GPO) build a prototype for a national packet switching network, Davies emphasized that:

> Such an experiment at an early stage is needed to develop the knowledge of these systems in the GPO and the British computer and communications industry ... It is very important not to find ourselves forced to buy computers and software for these systems from [the] USA. We could, by starting early enough, develop export markets.

Davies's concern with issues of economics and user-friendliness underscores the national context in which he conceived the idea of a packet switching network. Davies did not envision a world in which his proposed network would be the only surviving communications system. Rather, he saw a world in which packet switching networks would need to compete with other communications systems to attract and serve the business user, and a world in which Britain would need to compete with the United States and other countries to offer innovative computer products.

In December 1965 Davies proposed the idea of a national packet switching network that would provide low-cost data communications across Britain. He envisioned the network as providing a number of services to business and recreational users, including remote data processing, point-of-sale transactions, database queries, remote control of machines, and even on-line betting. In his scheme, a backbone of high capacity telephone lines would link major cities in the United Kingdom; the proposed network had multiple connections to most nodes, although it was not nearly as redundant as Baran's system. Like Baran, however, the NPL group designed the network with a dynamic, distributed routeing system, with each node making routeing decisions independently, according to current conditions in the network. The nodes would be connected by high-speed telephone lines so as to provide fast response times for interactive computing. Users would attach their computers, terminals, and printers to the nodes through dedicated interface computers at local sites.

Davies was convinced that the type of data communications infrastructure he was proposing would be necessary to keep Britain competitive in the information age. However, NPL did not have the resources or authority to build such a large network on its own. This authority belonged to the GPO which ran the national postal and telephone networks, but the managers there had little knowledge of or interest in data communications. In consequence, Davies was only able to build a small prototype network, called the Mark 1, rather than the nationwide system he had proposed….

PUTTING IT ALL TOGETHER: PACKET SWITCHING AND THE ARPANET

Baran and Davies had both envisioned nationwide networks that would use the new technique of packet switching, but neither one had been able to fully realize this goal. Instead, the first large-scale packet switching network would be built by the Advanced Research Projects Agency (ARPA). The design of this network would draw on the work of both Baran and Davies, but the network's builders had their own vision of what packet switching could achieve.

ARPA was one of many new American science and technology ventures that had been prompted by the cold war. Founded in 1958 in response to the Soviet Sputnik launch, ARPA's mission is to keep the US ahead of its military rivals by pursuing research projects that promise a significant advance in defense-related fields. ARPA is a small agency with no laboratories of its own; ARPA staff initiate and manage projects, but the actual research and development is done by academic and industry contractors. The agency is recognized even by critics for its good management and rapid development of new technologies, and has had some notable successes in transferring its technologies to the armed services and the private sector.

In 1962 ARPA…became a major funder of computer science in the United States, often exceeding university funding by significant amounts. IPTO (Information Processing Techniques Office [one of the project offices founded in 1962]) has been the driving force behind several important areas of computing research in the United States, including graphics, artificial intelligence, time-sharing operating systems, and networking.

ARPA's funding of basic research fitted in with the philosophy of Lyndon Johnson's administration. President Johnson advocated the use of agency funds to support basic research in universities in a September 1965 memo to his cabinet. Noting that funding by various federal agencies made up about two-thirds of total university research spending, he said that this money should be used to establish 'creative centers of excellence' throughout the nation. He urged each government agency engaged in research to take 'all practical measures ... to strengthen the institutions where research now goes on, and to help additional institutions to become more effective centers for teaching and research'. Johnson specifically did not want to limit research at these centres to mission-oriented projects: 'Under this policy more support will be provided under terms which give the university and the investigator wider scope for inquiry, as contrasted with highly specific, narrowly defined projects.

A few months later the Department of Defense responded to Johnson's call with a plan to create 'centers of excellence' in defence-related research. According to this plan, 'Each new university

program should present a stimulating challenge to faculty and students and, at the same time, contribute to basic knowledge needed for solving problems in national defense'. IPTO created several computing research centres, giving large grants to universities such as MIT, Carnegie Mellon, and UCLA. By the end of the decade ARPA had funded a variety of time-sharing computers located at universities and other computing research sites across the United States.
The purpose of its proposed network - the ARPANET - was to connect these scattered computing sites.

The ARPANET project was managed by Lawrence Roberts, a computer scientist who had conducted networking experiments at MIT's Lincoln Laboratories before joining ARPA in 1966. Roberts had a mandate to build a large-scale, multi-computer network, but he did not initially have a firm idea of how this should be done. He considered having pairs of computers establish a connection using ordinary telephone calls whenever they needed to exchange data, a method he had employed in his earlier experiments. But the high cost of long-distance telephone connections made this option seem prohibitively expensive. Roberts also worried that the ordinary phone service would be unacceptably prone to transmission errors and line failures. Although he was aware of the concept of packet switching, Roberts was not sure how to implement it in a large network.

With these issues still unresolved, Roberts attended a computing symposium in Gatlinburg, Tennessee, in October 1967, where he was to present ARPA's tentative networking plans. Roger Scantlebury from NPL also presented a paper, and Roberts heard for the first time about Davies's packet switching ideas and the ongoing work on the Mark I. Afterwards, a number of conference attendees gathered to discuss network design informally, and Scantlebury and his colleagues advocated packet switching as a solution to Roberts's concerns about line efficiency. The NPL group influenced a number of American computer scientists in favour of the new technique, and they adopted Davies's term 'packet switching' to refer to this type of network. Roberts also adopted some specific aspects of the NPL design. For instance, Roberts had planned to use relative low-speed telephone lines to connect the network nodes. He later recalled that after the NPL representatives 'spent all night with me arguing about the thing back and forth . . . I concluded from those arguments that wider bandwidths would be useful'. The ARPANET also used a packet format similar to NPL's. After the ARPANET project was underway, the firm of Bolt Beranek and Newman which was awarded the main contract to build the network nodes, continued to interact with the NPL group.

Baran also became directly involved in the early stages of planning the ARPANET. Scantlebury had referred Roberts to Baran's earlier work, and soon after his return from Gatlinburg Roberts read Baran's *On Distributed Communications.* He would later describe this as a kind of revelation: 'suddenly I learned how to route packets'. Some of the ARPANET contractors were already aware of Baran's work and had used it in their own research. Roberts recruited Baran in 1967 to advise the ARPANET planning group about distributed communications and packet switching.

Through these various encounters, Roberts and other members of the ARPANET group were exposed to the ideas and techniques of both Baran and Davies, and became convinced that packet switching and distributed networking would be both feasible and desirable for the

ARPANET. The new technique promised to make more efficient use of the network's long-distance communications links as well as enhance the system's ability to recover from the equipment failures that an experimental network would surely encounter. At the same time, however, packet switching was still an unproven technique that would be difficult to implement successfully. The decision to employ packet switching on such a large scale reflected ARPA's commitment to high-risk research: if it worked, the pay-off would be not only greater efficiency and ruggedness in the ARPANET itself, but also a significant advance in computer scientists' understanding of network properties and techniques. ARPA managers could afford - indeed, had a mandate - to think extravagantly, to aim for the highest pay-off rather than the safest investment.

THE SOCIAL CONSTRUCTION OF PACKET SWITCHING

The three projects sponsored by Rand, NPL, and ARPA had much in common in their approach to packet switching, but also some crucial differences that helped the ARPANET play a more enduring and influential role than the others. Davies, Baran and Roberts each made technical choices based on specific local concerns, and the extent to which their systems were influential depended in part on whether others shared those concerns. For instance, Baran's system had many elements that were specifically adapted to the cold war threat, including the very high levels of redundancy, the location of nodes away from population centres, and the integration of cryptographic capabilities and priority/precedence features into the system design. None of these features was adopted by Davies or Roberts, neither of whom was concerned with survivability. On the other hand, aspects of Baran's system that would be useful in a variety of situations - such as high-speed transmission, adaptive routeing and efficient packet switching – did find a home in later systems....

…In both the United States and the United Kingdom, computing technologies became policy instruments in the 1960s. In Britain, intervention in the computer industry was seen as a symbol of the Labour Party's commitment to modernization as well as an engine of economic growth, and the government made efforts to fund research and coordinate industrial production. In the United States, technological prowess was seen as a weapon in the cold war, and defence-related research was generously funded through organizations like Rand and ARPA. In both countries, individuals and organizations interested in pursuing computer networking often found it necessary to join government-sponsored projects or to present their work as responsive to contemporary political agendas….

The fact that packet switching had to be integrated into local practices and concerns led to very different outcomes in the three network projects. Some visions of packet networking were easier to implement, some turned out to be a better match for evolving computer technology, and some were more attractive to organizations in a position to sponsor network projects. Making packet switching work was not just a matter of having the right technical idea: it also required the right environment. Only after the ARPANET presented a highly visible example of a successful packet switching system did it come to be seen as a self-evidently superior technique. The success of the ARPANET may have depended on packet switching, but it could equally be argued that the success of packet switching depended on the ARPANET.