Introduction to Davies' paper

In early March of 2000, my morning electronic-mail unexpectedly contained a paper from Donald Davies which is printed below. Its opening sentence makes clear his motivation: '…to help those who study the early history of computer networking'. He amplified his motivation in a separate private electronic note: 'Part of the reason for writing it was to guide [a colleague at] CERN who is in the last stages of writing an excellent book…and was becoming puzzled by what he read [about the genesis of computer networking]'. In yet a different note, he also voiced the feeling that 'this will be the last "historical document" that I will write for a while'. Finally, he expressed a wish not to have the paper published at the time; his health could not have supported a public debate. Comments in other messages suggested a premonition that his medical condition would deteriorate; and it did with his death on May 28 of 2000.

It is now a year after the final draft of his historical contribution was completed and privately circulated. It is not certain to whom or how many copies were distributed, but about a dozen are known. There may be others. It is appropriate to honor Donald Davies' intention that this contribution be posthumously published for the benefit of historians and other scholars who study the progress of computer networking. It also seems proper, given Donald Davies' Fellowship in the Royal Society and his association with the National Physical Laboratory and other United Kingdom organizations, that the publication be in a British journal. Accordingly, it was submitted to the British Computer Society for consideration and has been graciously accepted for its Journal.

Donald Davies' paper is a very thorough, even exhaustive and certainly scholarly, analysis of several documents that have been cited and publicized in the media and on Web pages as having established a priority claim for invention of the construct now known as packet-switching. Davies was very troubled by what he (and others) believed was an improper action. In this paper and with a carefully balanced argument, he directs attention to what he regarded was the use of misrepresented documents to support such a claim.

His discussion of the referenced documents in question—some of them heavily mathematical—is augmented by three appendices and careful citations of relevant publications. His argument is persuasive and hopefully will resolve the conflicting claims about the origin of the construct of packet-switching, the origin of the name itself and the contribution and role of other groups and organizations during the earliest days of the Net.

Willis H. Ware
Santa Monica, California, USA

Biographical Note

Dr Willis H. Ware, trained as an electrical engineer, is a senior computer scientist emeritus with the RAND Corporation (1952–present) in Santa Monica, California where he has held various staff and management positions. His research interests in the security of information systems began in the late 1960s and now focus on the technical and policy aspects of highly automated and computerized information-oriented societies. He continues to be active professionally as speaker, conferee and researcher, and is the recipient of numerous honors and awards.

An Historical Study of the Beginnings of Packet Switching

D. W. DAVIES

Deceased

1. INTRODUCTION

The purpose of this paper is to help those who study the early history of computer networking by providing pointers to primary documents concerning one of the first steps in the new technology—the introduction of packet switching. I will give my own interpretation of these documents, as one of those who took part in the developments, and this interpretation can be checked for accuracy from the documents themselves.

This paper has three annexes which are representative of the early work on packet switched systems at the UK National Physical Laboratory (NPL). Annex 1 is the first note on the subject that I wrote on November 10th 1965. Annex 2 contains part of a note dated December 15th 1965. Only those passages with new material are in this transcript. Annex 3 is a brief summary of a paper [1] by the NPL team presented at a symposium in Gatlinburg in October 1967.

There is a good degree of consensus about this early history among those who have written at length in well-researched books [2, 3, 4]. This asserts that packet switching had two independent beginnings, with Paul Baran and Donald Davies, the present author.

Briefly, the first description of packet switching is attributed to Paul Baran [5] who proposed it in the context of a digital communication network for military use, where a prime requirement was an ability to continue working after considerable damage to the switching centres and the links between them, without losing data if at all possible. The digital data would include digitized speech. This short summary does not do justice to the extensive study by Paul Baran and his team at RAND Corporation. The full text of the report can be found on the Web [6].

Independently, starting from a different motivation, the author of this paper arrived at a striking similar digital communication system. My first three notes for private circulation were written in late 1965 and a more complete report was written in June 1966 and widely circulated. (I hope that these papers will soon be scanned and made available on the net.) The first open publication [1] was at an ACM seminar in Gatlinburg in October 1967. Wider publication of the early work we did at NPL was made at the 1968 IFIP Congress [7, 8, 9, 10] and a serious historical study of the whole of our data communications work, by Martin Campbell-Kelly, was published in 1988 [11].

Packet switching was adopted for the ARPA network after the 1967 symposium, but it remains uncertain whether it was known to Lawrence G. Roberts before that meeting. He gave a paper [12] on the early ARPA plans in the same session at Gatlinburg, which is an important primary document bearing on that question. There have been references to a report giving very early ARPA network designs, which could perhaps be sought from US Government sources.

Two Web pages [13, 14] by early workers in the field give alternative accounts of the beginnings of packet switching, with a different chronology. It is the resolution of these differences which is the purpose of this paper.

2. DIFFERENT VIEWPOINTS

Leonard Kleinrock heads the Web page [13] 'Brief summary of firsts, key accomplishments and contributions from Len Kleinrock' and items 1, 3 and 4 of his 16-item list make a definite claim to be the originator of packet switching.

Leonard Kleinrock made a very significant contribution to the development of the ARPA network by the skilful application of queuing theory to predict performance and by studies of the optimization of its design. He is particularly noted for a book [15] which applied queuing theory to a general design of store-and-forward networks.

Lawrence Roberts was the driving force behind the ARPA network development and, together with his team, must be credited with building the first packet-switched wide area network, thus demonstrating beyond doubt that the new concept had made a great advance in communications technology. His Web page [14] is headed 'Internet Chronology' and his items 1 and 4 of an extensive list support Kleinrock's claims.

Roberts quotes an earlier paper of Kleinrock [16] as the 'first paper on packet switching theory'. At about the same time, May 1961, Kleinrock submitted a proposal [17] for a Ph.D. thesis 'Information flow in large communication nets' which was accepted by MIT. The resultant thesis was the basis for his book [15] which, it can be assumed, contains the distilled wisdom of the earlier publications. Roberts writes of [16]: 'This was the theoretical work that convinced Roberts that packets could be used for the Internet.'

Roberts then quotes [15] as providing 'the network design and queuing theory necessary to build packet networks'. 'This work was a major factor in designing the communications network for the ARPANET. It shows that packet switching would work, whereas until ARPANET was built in 1969, most communications experts claimed that packet switching would never work.'

By writing this paper I hope to guide historians who want to resolve the differences between these views of Roberts and Kleinrock and the others who have written about this early step in network technology.

3. PACKET SWITCHING AS A STEP FORWARD IN COMMUNICATIONS TECHNOLOGY

Packet switching seems to me a distinct step in our understanding of communications technology. After the ARPA work, the Internet and the Web were based on this technology but, to some extent, their success inhibited development for a short period.

Now the introduction of 'asynchronous time-division multiplexing' or ATM by the ITU, which sets international standards, is set to bring the technology of packet switching into line with other technical developments, so that it may become the common method of multiplexing and switching for all telecommunications, including speech. Roberts' chronology, updated to 1997, moves the picture forward to 1998, 2000 and 2005 to predict the influence of ATM. This perhaps makes it urgent to understand just how packet switching began.

I believe it may happen that confusion about terminology leads to misunderstanding. Sometimes new technical terms are needed in order to be able to speak clearly about a new concept. In 1966 I saw the need to distinguish the unit of information in which a communication subnet operated (the packet) from the unit known to its end users. After discussion with a linguist working on our Russian to English machine translation project, Steve Whelan, I decided on the word 'packet'. To be quite certain we are using the same language I will divert here to restate what, to most people, must be obvious.

4. MESSAGE SWITCHING AND PACKET SWITCHING

There should be no problem in making the distinction between message switching and packet switching but, in case this has been a cause of misunderstanding in the claims of Kleinrock and Roberts, I will identify the differences.

4.1. Message switching

When the 1964 book *Communication Nets* by Kleinrock was published, message switching was an established technique. An example was used in Chapter 1, Section 1.5 of that book, entitled 'An existing store-and-forward communication net', to illustrate the kind of system that his research would analyse. His description was condensed from a paper in the Western Union Technical Review and covered their 'Plan 55A' for use by the Air Force. Telegraph messages were sent between switching centres and received on punched paper tape. Routing indicators in the messages showed where each message should go next. At the switching centre the contents of a tape were transmitted internally to a 'sending station' for its next centre and a new tape was punched. These tapes formed a buffer or queue for the messages going out on this line. When their turn came they were transmitted again.

Recording on tape and resending twice at a centre is perhaps unusual. Other variants of these message switching centres have the paper tape carried physically to the sending station. The general principle became known as a 'torn-tape' switching centre.

At some stage, as computer technology improved, messages in these centres were stored on a magnetic drum and made accessible to the sending mechanism from the drum.

The lines between centres ran at standard telegraph speed, typically 50 baud. If the traffic demanded more capacity, extra low-speed lines were employed. Evidently the queuing time was not regarded as a critical parameter or higher line speeds would be used. Delays across the network in hours were accepted.

Given a choice of interoffice line speed, the transmission time was proportional to message length. This was at the choice of the message originators, so the transit delay distribution was not controlled by the network operator—a severe drawback if interactive use of computers was one purpose of the network. When a switch had received a message it was fully responsible for its safety. In the absence of any end-to-end procedures, messages could not be re-sent if they were lost. For example, a received message was stored fully on the drum and its integrity checked before doing anything with it and sometimes more than one drum was employed for safety.

Kleinrock's analysis of networks similar to his example has to employ a realistic model for rates of message origination and message lengths. The simplest models, well known for their tractability, use message origination with a Poisson model and message length exponentially distributed, and these are what Kleinrock employed. He needed to make a further statistical simplification (the independence assumption) to obtain a fully tractable model.

Exponential message-length distribution gives no upper limit on message length, which is perhaps an extreme assumption and is quite unlike what we would expect with packet switching. Interestingly, when Kleinrock began to analyse the ARPA net he retained the exponential length distribution, but with a much smaller average length to reflect the new design feature.

4.2. Packet switching

Paul Baran used the essential principle of packet switching by breaking the messages received by his network into 'short message blocks' of size 1024 bits. An end-to-end procedure arranged for the packets received at the destination to be reassembled into their original message. This had several virtues, principally that the transit delay was reduced and was under the control of the network operator. It reduced the need for memory at the switching centres and it often allowed immediate recovery from a failure of a centre or a line, so that following packets could be rerouted.

Baran's network also had to carry digitized voice traffic, where the transit delay is very critical. This requirement certainly demanded packets.

My own work was aimed at providing a general purpose communication network for time-sharing computers and their users, very much like what the ARPA net became and perhaps even closer to the present form of the Internet, since commercial use was part of the scheme. Response time was of paramount importance to get a truly interactive 'feel' for the services provided. I regarded 1 s as an absolute limit and aimed for 0.1 s if possible. My first rough application of simple queuing theory indicated that 0.1 s was possible—to my astonishment—since I had the received idea that store-and-forward must be slow.

In conclusion, the essential feature of packet switching as an advance over message switching was the splitting of submitted messages into packets with a strict limit on their size. The purpose was to obtain a smaller transit delay which is under the control of the network operator.

Before leaving this subject, may I remark, as Kleinrock has done in email correspondence with me, that successful implementation of packet switching requires much more than just the use of this relatively simple principle. Indeed the principle itself proves complex when protocols have to be developed, such as TCP/IP, to govern the operation of the principle and inter-working of divergent packet networks is also needed. Routing of packets had been studied by Baran and needed more development. Flow control was needed as well as a mechanism to choke off demand in places and prevent destructive congestion. The ARPA network project would find and demonstrate solutions to most of these problems.

But, regardless of the complexities that networking would uncover, an essential first step was the recognition that a new scheme of communication was needed; that is to say, packet switching. This paper concerns that essential first step.

5. THE EARLY WORK OF LEONARD KLEINROCK

I have access to two documents, Kleinrock's proposal for a Ph.D. Thesis [17] and the 1964 book *Communication Nets—Stochastic Message Flow and Delay* [15] in the McGraw-Hill edition. The proposal [17], entitled *Information Flow in Large Communication Nets* is dated May 13th 1961 and comes from the MIT Research Laboratory of Electronics (RLE). It was approved on July 24th 1961.

Roberts in his Web page quoted here refers to a 'first paper on packet switching theory' with the same title as the thesis proposal, in a RLE quarterly progress report [16] of July 1961. This is evidently Quarterly Progress Report No. 62 dated July 15th 1961, pages 162–163, and must be a brief summary. It is the single reference [2] in an apparent addendum to the thesis proposal labelled D, Information Flow in Large Communication Nets which begins: 'In our continuing research on the problems of information flow in large communication nets results have been obtained (for a single node) for two classes of queue disciplines: priority queuing and time-shared servicing....' The material need not be considered as a separate issue because it is largely repeated in a more elaborate form in Chapter 5 of the book. In particular, Section 5.3 on Time-shared Service in pages 84–94 requires special attention later.

The 34 pages of the proposal, including a bibliography of 30 items, thoroughly covers the existing literature before presenting new results and extensions of previous work in the form of five theorems. A 12-page appendix contains proofs of the theorems. I could find no reference in the proposal to packet switching in the sense defined above, namely the splitting of incoming messages into smaller units or packets and their reassembly at their destination.

The 1964 book has been said by Kleinrock to contain the essence of the thesis results. By using the book to evaluate his contribution to the concept of packet switching we shall not do him an injustice, because clearly any new work after the thesis which he considered significant for the understanding of network behaviour will have been incorporated in the book. In particular, if the work in [16] had been a 'first paper on packet switching theory' as Roberts asserts, this would surely have been incorporated in the book three years later. If Kleinrock had, for some reason, not recognized the importance of this supposed 1961 invention of packet switching, so that it did not appear in the book, this would cast doubt on his claim to priority.

Writers wishing to get close to the true story are recommended strongly to study very carefully the contents of *Communication Nets*, the 1964 book. It is my contention that if packet switching has not been clearly signalled in that book, and its importance recognized, its invention cannot be credited to Leonard Kleinrock. I will give the results of my study of the book but others must form their own judgements.

The preface to the book outlines its scope, which is to consider and analyse the flow of message traffic in connected networks of communication centres. The chief measure of performance is the average delay encountered by a message in passing through the net. The contents of the seven main chapters are summarized in the preface and the eighth chapter proposes further investigations. These are followed by proofs of the theorems used in the book. The book is said to be based on the doctoral thesis. As promised in the thesis proposal, design optimization is an important theme and routing principles are covered in Chapter 7, while a simulation study is also contained in that chapter.

Basic ideas are given at the start of Chapter 1, followed by a specific example of a message communications network in Section 1.5. This is the Western Union message-switching system employing the form of torn-tape switching centre which I described above as a prime example of the existing state of message switching before the invention of packet switching. To clarify the absence of packet switching in this example network, note that its message traffic, specified in Section 1.6 'the model assumptions', has message lengths which are exponentially distributed. This model, with no upper limit on message length, is typical of what might be expected with message traffic instead of packets of substantially constant size, generated by splitting the given messages into packets. Ironically, constant sized packets would have made a tractable model for analysis. In the sections on 'Notation and Further Definitions' and also 'Precise Problem Statement', the choice of exponential message length distribution is implied, but confirmation depends on the results summarized in Chapter 2, where the average message length $1/\mu$ is applied to all traffic. We note that a further assumption for the statistics of message traffic was found necessary to have a tractable analysis, the so-called 'independence assumption', according to which traffic arriving at each node's output queue has individual message lengths chosen afresh from an exponentially distributed population with average length $1/\mu$.

I have to report that nothing in the book seems to envisage packet switching. If the ideas developed by Paul Baran or me were present, I would expect to recognize them easily, and I do not. However there is the possibility that they may be hidden in some way. For this reason I asked Leonard Kleinrock to identify the places where packet switching was treated. His reply was to indicate Section 5.3 of the book. This is closely related to reference [1] in his thesis proposal (reference [17] here), and is quoted by Roberts as the theoretical work that convinced him that packets could be used for the Internet. It therefore demands our special attention.

This important section can be looked at in two ways. Firstly we could take it on its merits, reading no more than is written and regarding its conclusions literally. Secondly we could see where its thoughts lead us and hypothesize that Kleinrock may have followed the same line of thinking. This second approach carries the risk that we may be making the invention of packet switching on Kleinrock's behalf.

5.1. A literal view of Section 5.3

The chapter title is 'Waiting times for certain queue disciplines'. It aims to explore 'the manner in which message delay is affected when one introduces a priority structure (or queue discipline) into the set of messages in a single node facility with a single transmission (or service) channel'. This is a modest target and tells us little about the net as a whole, but perhaps an understanding of the simple case will throw light on the whole network problem?

In Section 5.1, a set of four queue disciplines is analysed. The expected time spent in the queue is plotted for each case. An interesting conservation law is given in theorem 5.4, proved in an appendix. This law provides an important constraint on the behaviour of any plausible priority scheme.

Section 5.3 seems to be a new theme, entitled 'Time-shared Service'. It presents results for a simple 'round-robin' time-shared service facility. It is said that 'Such a scheme is a likely candidate for the discipline of a large time-shared computational facility.' If this were the only outcome, it would have no relevance to packet switching.

Time is quantized into segments of length Q. Items in the queue are serviced for this time, then the remaining part of the message goes to the back of the queue and the next message has its run of length Q. Each message recycles until there have been enough periods of length Q to exhaust it. The whole message is serviced no faster than if it had received the full use of the channel, but on the other hand, messages do not have to wait their turn but receive some service very rapidly. I have here ignored some fine detail in the service scheme, which has little effect on the results.

It is easy to understand what this implies in a single queue situation. It is also clear how it applies to a time-shared computing service, which is its obvious application. What is not clear is the treatment of the partial messages which result from the time-shared service. I think we have to assume that they arrive in sequence at their next node and become a message again, for treatment according to the queue discipline there. Following through such a scheme would have made an interesting study, but Kleinrock does not go beyond the single queue.

Section 5.4 extends the results for priority queuing to a more complex case in which priority increases with the amount of delay already experienced. This is unrelated to the time-sharing scheme.

As an overall result, it is stated that 'a meaningful average of the average waiting times is invariant with respect to a change in queue discipline.' Taken literally, this would mean that packet switching had no effect on end-to-end delay. Nothing is written which encourages the idea of reducing delay by splitting messages into packets.

5.2. A liberal view of Section 5.3

Having decided to investigate a queue discipline in which messages are treated in steps, in a round-robin manner, Kleinrock could have considered these short pieces of a message as individual entities. The advantage in reduced delay, especially when the pieces are very small, would have been obvious. There would be two approaches.

One would be to analyse message switching in which the messages were indeed these short pieces, or packets, of constant length. The individual discipline, M/D/1, is well known. This is the approach I used in my earliest papers. The statistical approximation involved was hard to justify, but when an extensive network of this kind was simulated a few years later, the results were gratifyingly close to the simple theory. The approximations used in Kleinrock's main work were perhaps more firmly based. But my aim was to get results which could be tested by later experiments, rather than to obtain mathematically defensible formulae.

A second approach for Kleinrock would have been a simulation of the packet-switched system in sufficient detail to get rough, practical results for end-to-end delay. Although he reported extensive simulation work in Chapter 7 of the book, nothing of this kind appears.

It may be relevant for me to discuss the big difference in method which I observe, as between Paul Baran and myself on the one hand and Leonard Kleinrock on the other. I surmise that Baran's aim was to develop a workable design for a survivable network. This did not necessarily mean very accurate analysis of performance, only such analysis as was needed to understand behaviour and get rough performance data. Certainly this was my method of working as, for instance, in my very rough use of queuing theory to predict the performance of packet switching—just one formula sufficed.

On the other hand, it seems that Kleinrock's motivation was different. His analysis of complex systems was masterly and went well beyond what had been achieved before. But basically he was interested in systems that he could analyse with his customary skill. Getting a very rough result for a different type of system did not usually attract him. My conclusion is that this was the reason that Kleinrock did not make the leap of thought needed to arrive at packet switching. In fact I suspect he would have regarded it as a rather trivial extension to his work.

5.3. Conclusions concerning Kleinrock's work prior to 1964

In my view, the principles of packet switching are fairly simple and easily stated. Both Paul Baran's work and mine contain such statements. In both our cases the description and analysis in our papers is thorough and complex, but a simpler description is possible and would easily be recognized.

In Leonard Kleinrock's work, as revealed in the 1964 book and the thesis proposal, I can find no evidence that he understood the principles of packet switching. Because of the very significant effect of packet switching on the transit delay of messages, any discoverer of this technique would be expected to explain and demonstrate this effect in detail, either by mathematical analysis or simulation. Kleinrock does neither. His one possible claim is based on Section 5.3 of this book. I have concluded that by following through the ideas in that section, he could have arrived at the concept of packet switching, but that he did not do so.

If it is true that Kleinrock had made no contribution to the theory of packet switching in 1964, then it is difficult to understand the statements in Lawrence Roberts' *Internet Chronology*. How did the July 1961 paper 'convince Roberts that packets could be used for the Internet' when packets appear nowhere in Kleinrock's work before 1964?

The most favourable conclusion I can make is that Kleinrock's book was a factor in convincing Roberts that message switching would work. In this sense it could help him to convince sceptical backers that there was a mathematical theory covering the queues which appeared in such a network. The distinction between packets and messages could well be beyond their understanding.

There can be no doubt that around 1967, at the time of the Gatlinburg symposium, Roberts understood the implications of packet switching. His statements about what he knew from Baran and myself at the time have not always been self-consistent and it is not possible to know just when he understood the situation. Certainly, there is absolutely nothing in his Gatlinburg paper [12] to show that the ARPA design at that time took account of packet switching or of even the simplest traffic considerations.

6. SUBSEQUENT HISTORICAL PAPERS

6.1. 1970

At a session of the Spring Joint Computer Conference in Atlantic City, May 7th 1970, five important technical papers were read, including one each by Roberts and Kleinrock. A study of these papers may help the historian to understand how these two participants saw the history of packet switching at that time. ARPA recognized the importance of this session by reprinting the five papers in its own binding.

The introductory paper was 'Computer network development to achieve resource sharing' by Lawrence G. Roberts and Barry D. Wessler [18]. This account of the design of the ARPA network references the work of Baran as a 'fully distributed message switching system'. It does not reference Kleinrock's book or his 1961 paper, which is strange in the light of Roberts' later attribution of confidence in packet switching to Kleinrock's work. The acknowledgement of Kleinrock's work in Roberts' paper is to his paper in the same session [19] and there is also a reference to simulation work by Frank *et al.* [20] who formed Network Analysis Corporation and were under contract to ARPA for optimization and other work.

Kleinrock's paper in the same session [19] contains a summary of theoretical and simulation work bearing on the design of the ARPA network, essentially a development of the work described in the 1964 book, with more on routing algorithms. It shows good agreement between analysis and simulation for 'small packet traffic only' and promises to extend this to multipacket traffic. There are no claims to have originated the concept of packet switching.

6.2. 1978

The session in 1970 throws an interesting light on the viewpoints of Roberts and Kleinrock soon after the start of network development. A set of papers in the Proceedings of the IEEE in 1978 is of even greater interest because the writers were evidently conscious of the historical content of their papers.

The November issue of *Proceedings of the IEEE* was devoted to packet communication networks, with an introduction [21] by Robert E. Kahn, Keith W. Uncapher and Harry L. Van Trees and papers by Roberts [22] and Kleinrock [23], among others. The whole volume was given a hardback binding and gold-lettered copies were presented to authors and others, including Paul Baran and myself.

The introductory paper clearly and simply ascribes the first presentation of packet switching to Paul Baran.

The paper by Roberts is probably the best historical review of packet switching that has ever been written, covering not only ARPA but other less well known but important projects such as SITA, TYMNET, RCP and EIN.

Issues now largely forgotten, such as 'Datagrams' versus 'Virtual Circuits' are discussed at length.

The sequence of early development according to Roberts is unambiguous. First came the RAND reports, describing a superior system but largely ignored until packet switching was rediscovered and applied by others.

Second in Roberts' chronology was 'ARPA I' which was the first concept of the resource sharing network. Licklider provided the intellectual support for the idea and the ARPA backing, but a network project had not yet begun.

Next, Roberts describes the early NPL work, our 1965 and 1966 papers and our 'discovery' of the RAND work.

The Gatlinburg symposium and our paper's proposal was 'similar to the actual networks being built today'. The cost analysis showed 'strong economic advantages for the packet approach' but our practical work was confined to a local network. However, this project 'plus the strong conviction and continued effort of those at NPL (Davies, Barber, Scantlebury, Wilkinson and Bartlett) did gradually have an effect on the UK and much of Europe'.

The next and the major item in his chronology is 'ARPA II' which is summarized, leading up to the major public demonstration at ICCC Washington 1972. It is in this connection that Kleinrock's work is mentioned but with no references to Kleinrock's work before 1972.

It is clear that Roberts in 1978 did not attribute any original work on packet switching to Leonard Kleinrock, a marked difference from his current *Internet Chronology*.

Kleinrock's paper in these proceedings [23] gives a good and clear account of topological considerations in network design, including topics such as deadlocks, degradations, distributed control, flow control and protocols.

In the historical part of the paper, Kleinrock refers the reader to the more comprehensive review by Roberts [22] but not before mentioning specifically the work of Paul Baran and the NPL team. He makes no claim to have originated packet switching in any form.

7. CONCLUSIONS

I can state here my own conclusions from a reading of the original papers which are listed below. These may be biased by my own participation in the early work, so it must be the task of the historian who studies these papers and others to arrive at an authoritative conclusion.

My contention is that the work of Kleinrock before and up to 1964 gives him no claim to have originated packet switching, the honour for which must go to Paul Baran. The passage in his book on time-sharing queue discipline, if pursued to a conclusion, might have led him to packet switching, but it did not.

Roberts must surely be wrong to attribute packet switching to Kleinrock, but the reason for this aberration remains unclear. Nothing in the 1961 paper he refers to seems relevant, but our knowledge of that paper is incomplete. His reference to the 1964 book may arise from the fact that the presence of an impressive body of queuing theory for message switching could impress those who backed the ARPA project, with no clear distinction between packet and message networks. What Roberts learnt from the NPL paper at the Gatlinburg symposium must remain uncertain. I am convinced he should have learnt a lot from that comprehensive paper.

The Web pages by Kleinrock [13] and Roberts [14] are in my opinion very misleading. Any person studying the early history of networking is urged not to take them at face value but to study the original documents in order to get a more accurate picture.

REFERENCES

[1] Bartlett, K., Davies, D., Scantlebury, S. and Wilkinson, P. (1967) A digital communication network for computers giving rapid response to remote terminals. *ACM Symp. Operating Syst. Principles*, October. ACM Press.

[2] Abbate, J. (1999) *Inventing the Internet*. MIT, Cambridge, MA.

[3] Hafner, K. and Lyon, M. (1996) *Where Wizards Stay up Late*. Simon and Schuster, New York.

[4] Naughton, J. (1999) *A Brief History of the Future*. Weidenfeld and Nicolson, London.

[5] Baran, P. (1964) On distributed communications networks. *IEEE Trans. Commun. Syst.*, March, 1–9.

[6] Baran, P., Boehm, S. and Smith, J. W. (1964) *On Distributed Communications*, RAND Corporation Memorandum RM-3420-PR. Available on the RAND site www.rand.org by following 'publications' then 'RAND classics'.

[7] Davies, D. W. (1968) The principles of a data communication network for computers and remote peripherals. In *Proc. IFIP Congress on Information Processing 68*. North-Holland, Amsterdam.

[8] Wilkinson, P. T. and Scantlebury, R. A. (1968) The control functions in a local data network. In *Proc. IFIP Congress on Information Processing 68*. North-Holland, Amsterdam.

[9] Bartlett, K. A. (1968) Transmission control in a local data network. In *Proc. IFIP Congress on Information Processing 68*. North-Holland, Amsterdam.

[10] Scantlebury, R. A., Wilkinson, P. T. and Bartlett, K. A. (1968) The design of a message switching centre for a digital communication network In *Proc. IFIP Congress on Information Processing 68*. North-Holland, Amsterdam.

[11] Campbell-Kelly, M. (1988) Data communications at the National Physical Laboratory (1965–1975). *Ann. History Comput.* **9**.

[12] Roberts, L. (1967) Multiple computer networks and inter-computer communication. *ACM Symp. Operating Syst. Principles*, October. ACM Press.

[13] Kleinrock, L. *Brief Summary of Firsts, Key Accomplishments and Contributions for Len Kleinrock*. www.lk.cs.ucla.edu/LK/Inet/keys

[14] Roberts, L. G. (1997) *Internet Chronology*. www.ziplink.net/ lroberts/InternetChronology

[15] Kleinrock, L. (1964) *Communication Nets, Stochastic Message Flow and Delay*. McGraw-Hill.

[16] Kleinrock, L. (1961) Information flow in large communication networks. *RLE Quart. Prog. Rep.*, **62**, 162–163.

[17] Kleinrock, L. (1961) *Information Flow in Large Communication Nets*. Proposal for a PhD Thesis, 1961.

[18] Roberts, L. G. and Wessler, B. D. (1970) Computer network development to achieve resource sharing. In *Proc. SJCC*, May 7th 1970, Atlantic City, pp. 543–549. Reprint by ARPA.

[19] Kleinrock, L. (1970) Analytic and simulation methods in computer network design. In *Proc. SJCC*, May 7th 1970, Atlantic City, pp. 569–579. Reprint by ARPA.

[20] Frank, H., Frisch, I. T. and Chou, W. (1970) Topological considerations in the design of the ATPA computer network. In *Proc. SJCC*, May 7th 1970, Atlantic City, pp. 581–587. Reprint by ARPA.

[21] Kahn, R. E., Uncapher, K. W. and Van Trees, H. L. (1978) Scanning the issue. *Special Issue on Packet Communication Networks*, *Proc. IEEE*, **66**, 1305–1305.

[22] Roberts, L. G. (1978) The evolution of packet switching. *Proc. IEEE*, **66**, 1307–1313.

[23] Kleinrock, L. (1978) Principles and lessons in packet communications. *Proc. IEEE*, **66**, 1320–1329.

ANNEXES

Annex 1 is a transcription from a very poor typescript containing the first note that I wrote on the subject. It was dated November 10th 1965. In Annex 2 a later note is quoted, in part, which was dated December 15th 1965. By March 1966 I considered my ideas for a new type of network sufficiently advanced to give a public lecture, which was exceptionally well attended. The content of that lecture was recorded in a paper dated June 1966 and widely distributed. I hope to have this scanned and made more widely available. Annex 3 is a brief summary of the paper by the NPL team presented at the 1967 Gatlinburg symposium. A paper by Lawrence Roberts on the new ARPA development was given in the same session.

Comments inserted into the transcript are in italic.

Annex 1. Remote on-line data processing and its communication needs

Remote on-line data processing is at a very early stage of practical achievement, so it is perhaps necessary to justify the making of long-term predictions, which are necessarily speculative. Long-term predictions may be valuable because they may indicate the future need for new kinds of communication service. The development of the telephone/telegraph is constrained by present-day economics and by the inertia to change due to the need for all equipment to fit into the existing system. If new services are needed, or better ways of providing the present services can be foreseen, requiring major changes, then planning well ahead is needed so that changes affecting the whole system can be started many years before they can be justified by immediate economics. It may also prove that research should be started now to determine how best to provide new services in, say, ten years' time.

This note contains some very tentative predictions that might go into a more prolonged and detailed study.

A1.1. Remote, on-line data processing

In this kind of data processing the computer system handles all the records in a well-developed file-store and communicates at about 10 characters per second with people working at keyboards with simple printers and, eventually, cheap tabular and other displays. (The displays must have local storage to keep the communication rate economically feasible.) Since data is not accumulated at stations, faster transmission using paper tape, cards or magnetic tape is not a requirement.

Larger users might have their own computers, which carry out simpler parts of the work, giving some intermediate feedback to the users, and collect the data so that it can be transmitted for major processing jobs. This would need a higher transmission rate.

If the use of individual stations becomes economic, and this depends on the efficiency of communication, this kind of traffic will be greater than that between satellite computer and central service, because many more small organizations will be able to afford the simple equipment.

Already the few such services in existence (such as computation, stock exchange and airline reservation) show that they can be economical. Improvements in the economics of computers seem to be continuing, therefore the expansion of such services will not be limited by the cost. The security of the information is a problem, however, which leads to a subject of research: Security of data in doing business via a national network. (This refers to security against unauthorized access rather than against loss of data.)

A1.2. Forecast of the traffic

A typical transaction will consist of about 50 characters sent at typewriting speed to the computing centre, and a reply of about the same size sent rather faster. These messages include identifying numbers and names, the request, confirmation of the request and reply.

The greatest traffic could only come if the public used this means for everyday purposes such as shopping. It is doubtful whether the public can work accurately and confidently enough in a fixed format and whether the £100 terminal cost (present-day prices allowing for mass production) would be acceptable. There may be limited use of tone-button dials for simple transactions, but this needs no addition to the telephone service to provide it. Ultimately, looking further ahead, speech recognition and natural language analysis might give the telephone service these new applications.

The big traffic for keyboard messages will be from professional operators. For example the staff dealing with the public in banks, Post Offices and Government offices will use them. Here the fixed format is not a difficulty. Other users will also be specialists in their job, such as designers using computational tools. People sending enquiries and placing orders of all kinds will make up a large section of the traffic.

Suppose, however, that 10% of the working population or 2% of the whole population, makes frequent use of these keyboards, sending 100 messages a day each. This might compare with non-business use of the telephone amounting to several telephone calls per head per day. Business use of the telephone might be reduced by the growth of the kind of service we contemplate.

The overall result is that telephone calls and short data messages will not be very different in number. Since any sensible engineering solution should pack a data message into the equivalent of a fraction of a second of a telephone channel, we can predict that the communication needs of data for on-line remote processing will be small compared with those of the telephone network.

We are assuming that the adoption of on-line data processing will largely remove the need for fast transmission of a considerable amount of data. Large quantities of data for transmission have either been accumulated over a period in the wrong place or generated by a computational process in the wrong place or generated by data acquisition machinery. The sort of situation where a large experiment generates data which must be transmitted for reduction elsewhere can be assumed to be exceptional, and not necessarily to require the most economical provision for communication.

A1.3. Qualities of the service needed

It is important to remark that, though its volume will be small, the short message traffic *(i.e. packet-switched)* will be vital to the country. It will, in fact, carry more information than the telephone network. Corresponding to its much lower redundancy, its requirement for noise-free transmission will be greater.

Each station sends and receives data at a low rate and sporadically. The unit message of about 50 characters probably takes 10 s to send to the centre and the reply takes 10 s to return. A delay of up to 1 s in each transmission might be permitted if it saved money.

When a station has connected with a certain data processing service and the service centre has identified the station, subsequent messages should be passed through the communication system without further red tape, but only occupying channel capacity when the keyboard is being used. Thus a keyboard could be permanently 'connected' to a distant service but only be occupying the minimum of local exchange equipment. The 'connection' would simply be an entry in a table in the controlling computer.

The accuracy worth striving for where humans generate and receive the data is limited, but one error in 10,000 characters would be a desirable goal. For those needing greater accuracy the system could provide for transmission by two paths and comparison, but the terminals would in that case need special attention in order not to add their own errors.

For security, the system might provide from a third source to each end of a transmission path a 'one time pad' on demand. This would ensure that nowhere in the network was the data in clear. There would, of course, be a limit to the degree of security that a public service could offer.

Accuracy of routeing is important and verification of the station's identities to each other is a service that will certainly be needed. This leads to a subject of research: the details of the interaction between the users and the system is such a message network (as distinct from the interaction between the two parties to a call).

(The section which follows was really the starting point of my investigations, the realization that store-and-forward delay could be minimized by using only short messages, i.e. packets.)

A1.4. Means of providing a short-message data service

There has been some experience with digital transmission for P.C.M. speech over junctions *(the UK name for short-distance trunks)*. This shows that, using regenerators at the spacing normally used for loading coils, two pairs of wires can carry 1.5×10^6 bps. The reduction of cost as compared with the transmission of digital data over a speech path is about 25:1 (Hartley and Thomas, IEE Colloquium, October 1965).

The channel capacity provided is embarrassingly large. If 10^6 people in the London area sit at keyboards for an eight-hour day and send 100 calls per day on average, each of 50 characters, this amounts to 1.4×10^6 bps on average *(during the working day)*. Thus the capacity of a single P.C.M. junction would not be used on any one route.

Since there must be independent channels to allow for fault immunity, it seems likely that digital data will share the transmission equipment provided for speech. It will not necessarily influence speech towards the universal adoption of P.C.M., in fact data might always have to be carried less efficiently than is technically feasible to suit the economics of speech transmission. Further speculation on the means of transmission is unnecessary since the cost of transmission can easily be seen to be very low.

The packing of the short messages *(packets)* onto a digital path at 10^6 bps is possible by short term storage with very small delays. For smaller capacities, down to about 10^4 bps there will be no problem, but since delays may occur at each switching centre they must be individually no more than 100 ms. This implies that 10^4 bps is about the lowest total transmission capacity between two switching centres which is allowable without partitioning the message into smaller pieces. Already at 50 characters the overhead due to routeing information will be relatively large and further partitioning seems undesirable.

A problem will arise in giving economical service in areas with a low concentration of data stations, such as residential areas. Apart from this, the cost of transmission will be very low, and it may happen that the capital cost of the computers handling the storage and multiplexing at each switching centre dominates the cost. This leads to a research topic: the system design and programming of message switching centres for the public network. Several examples of private systems exist but there is, I believe, no British firm in the business.

A1.5. Other uses for a short-message data service

Such a system could take over the telegraph and telex services and make them more convenient because the other applications would cover the cost of many more sets of terminal equipment. The interaction with the system might have to be simplified so that anyone used to the keyboard could send messages as well as operate his own specialized services. *(This is email by another name.)*

Ultimately the control of the switching system for telephones might use messages carried by the message network. This would be potentially cheaper and more reliable than using various forms of modulation on the telephone channels, and it would fit in with computer control *(of telephone switching)*. It would also allow more complicated interactions with the telephone system such as personal calls set up via keyboards.

10th November, 1965 D. W. Davies

Annex 2. Proposal for the development of a national communication service for on-line data processing

Starting from the assumption that on-line data processing will increase in importance, and that users of such services will be spread out over the country, it is easily seen that data transmission by a switched network such as the telephone network is not matched to the new communication needs that will be created.

The user of an on-line service wishes to be free to push keys sporadically, and at any rate he wishes, without occupying and wasting a communication channel. But he does not expect a reply from the computation service for less than a 'message' of several characters, typically between 10 and 100.

A message communication service in which short messages are temporarily stored in computers situated at the nodes of the network, and forwarded in turn, can give great economies in the use of transmission paths. Further economies are afforded by the use of digital transmission plant, with regenerators in place of linear amplifiers. The result of these two factors is that transmission cost can be extremely low by present-day standards.

A2.1. The terminal equipment

(At that time a teleprinter was the most likely terminal, since CRT display terminals were expensive. Uses such as transport, power utilities, banking and insurance, airline reservations, design offices and information services were mentioned. An exchange for 10,000 lines was envisaged, handling 200 messages per second. 100 such exchanges would generate about 10,000 non-local messages per second at peak and routes carrying 1000 messages per second would be needed.)

These very rough estimates are made in order to give numbers to assess the transmission and message handling problems. If terminals with displays become cheaper, this might encourage the more lavish use of the communication system and this, more than any other factor, could increase the traffic beyond the figures mentioned. For the case of printers only *(i.e. teleprinter terminals)* our figures are probably overestimates.

(By the time that the NPL local network was handling real traffic, display terminals had become the norm and the prediction of increased packet rates had come true.)

A2.2. Digital transmission

(The discussion in November is repeated, leading to the same conclusion—that the cost of transmission would be very low and could be ignored.)

A2.3. Tandem message exchanges

A node in the network which is not connected to terminals presents a much simpler computer problem than an exchange serving terminals, therefore it will be discussed first. Such an exchange *(i.e. the node)* can also be connected to several computer services. These services communicate at high speed in complete messages so they appear at the exchange much like the transmission paths of the network. They do, however, need a small extra amount of processing to add to each message a note of its source, and possibly to replace the stated destination by a route indication.

The task of such a tandem exchange is basically simple. Very simple ADT *(automatic data transfer)* devices could place the incoming messages into small buffers in a core store. A central computer could examine these buffers in turn and sort the messages, according to routeing information, into output queues. The output queues would fluctuate in length somewhat and waiting here would be the main source of delay. Simple ADT devices could transmit the messages at the rate allowed by the channels, whenever the queue was non-zero.

An approximate model (valid for the case of many channels) is that arrivals in the output queue are Poisson distributed and output is at a regular rate determined by the time t taken to transmit one message. We expect t to be about 1 ms. The mean time of waiting in the queue would then be

$$W = t\rho/2(1-\rho)$$

where ρ is the mean arrival rate divided by the steady (potential) departure rate, i.e. the degree of saturation. Since saturation would be disastrous to the service and traffic cannot be accurately assessed it is not likely that ρ would be greater than 0.8 and therefore the mean waiting time only $2t$.

This approximate model illustrates that appreciable queuing delays indicate near-disaster. If the arrival is less random than Poisson this is even more true.

Assuming that to carry a message no more than five tandem exchanges are ever needed, and therefore that such messages are held seven times in short buffers and five times in output queues and are transmitted six times, the total delay time would average about $23t$. This could be kept down to 100 ms if all the communication channels had a capacity of at least 250 messages per second. With digital transmission sort of capacity would easily be provided, and correspond to a few telephone (PCM) channels.

The rate of access to the store might however be a limitation, or additional core stacks a cost burden. But if a message fits into 10 store words and each message is handled three times in transit though a tandem exchange, a rate of 10,000 messages for the whole exchange would need only a single 3 ms store. Each queue would need two pointers which might also be in the core stores, but references to these in the main store might run a single store into access limitations for a 10,000 per second exchange.

For a rough estimate, the computer need would cost about £100,000 for the largest of such exchanges, allowing for considerable redundancy.

A2.4. Message exchanges with terminals

(This section describes a possible design for what became known as a PAD or packet assembler–disassembler, which would handle terminals with no processing ability, by the exchange of individual characters. In the eventual ARPA network this device came at a second stage of development as the TIMP or Terminal IMP, ready for the 1972 demonstration. A rough costing of the whole network showed the handling of terminals to be a major part of the cost. One suggested improvement was a terminal interface for the controlled exchange of characters, which in the NPL network was the 'British Standard Interface'.)

A2.5. External characteristics of the proposed message communication system

(These were (1) the use of ISO code, (2) terminals held a message until, for example, a 'new-line' character appeared to signal end of message, (3) maximum delay of 100 milliseconds and (4) high rate connections for attachment of computers.)

The communication system will therefore act as line concentrator, character assembler and message assembler for all the computers attached to it. This is obviously better than the present arrangements in which the telephone network does the line concentration and transmission and each computer does the rest for itself.

A2.6. Commands to the communication system

(This section was a discussion of the terminal to network and computer to network interfaces—a first attempt to define a protocol. But it contained nothing of subsequent importance.)

A2.7. Some uses for a message communication network

The original intention for its use, the connection of terminals to computer services, remains of primary importance. A selection of such services is listed:

- Numerical computation at various levels of generality
- Editing and typesetting of text
- Design services and problem oriented languages
- Availability of goods for sale
- Ordering of goods
- Invoicing, delivery notes, etc.
- Booking of transport
- Banking, establishing credit
- Remote access to national records, e.g. MPNI *(pensions)*, tax, police, medical, on a secure basis
- Betting.

The use of the system for people-to-people communication represents an elaboration of the Telex system. The greater distribution of terminals because of all the other services offered will make this method of communication *(email)* more useful.

The use of the system for machine-to-machine communication depends on its reliability, i.e. freedom from congestion and visible faults. It would be possible for vital data to be sent to the two nearest exchanges by separate conduits, for protection against pick and shovel. Examples of machine uses, many of them having very low data rates per station are:

- Road traffic control
- Monitoring and controlling of widespread plant, like pipelines, utility services or automatic meteorology stations
- Burglar alarms and other security devices connected to monitoring computers and giving regular reports when asked
- The control of the telephone switching system.

(This section concludes the paper with a brief discussion of bulk data transfer, the need for which should decrease.)

A2.8. The development of a message communication system

Proposal for a pilot service in London and for research and development in the UK. 'It is important not to find ourselves forced to buy computers and software for these systems from USA.'

15th December 1965 D. W. Davies

Annex 3. A digital communication network for computers giving rapid response at remote terminals

(D. W. Davies, K. A. Bartlett, R. A. Scantlebury and P. T. Wilkinson; summary by D. W. Davies)

Our paper at Gatlinburg had 10 pages and five diagrams. My summary will indicate the type of material included in each section.

Introduction. This paper proposes a design for a common-carrier data network. Reference is made to Baran's work and the current use of 1.5 Mb s^{-1} in short haul trunks and 224 Mb s^{-1} in experimental long-haul. 'The full economy of the proposed network may not be realized until a new type of local distribution network comes into operation', which presages developments now taking place!

Outline of the proposal. An outline description of a packet-switched network. Its 'high level network' consists of 'nodes' connected by digital links. These handle transit traffic and form the main network. Between these and users are 'interface computers' handling a variety of subscribers. Among these subscribers are the time-sharing computers which provide services. A response time of 100 ms is achievable. Resilience to failures is mentioned. Packet format is detailed in Figure 3.

The users of the network. Enquiry stations, file storage, bank proof-machines, printers, time-share services, remote sensors and actuators etc. In industry and commerce, sharing of data more important than access to remote computing. Control of transport and distribution, signalling system for the phone network.

Control procedures involving the users. A long section detailing interfaces for a tape reader and a keyboard with printer, as examples.

Design of the links and node computers in the high level network. A long section with a detailed design including special hardware to remove transmission chores from the node software. Figure 4 gives a logic diagram of the hardware, designed as an exercise for our local network.

Software organization of the node computer. Outline based on our software design for the local network by Peter Wilkinson.

Performance of the node. The node can handle 2600 packets per second. Delays were based on a queuing model with Poisson input and constant service time and independence between queues. This model includes queuing for processor time, based on the software model. Results in Figure 5. At 80% of saturation the mean delay per node was 1.5 ms. *(This proved optimistic concerning the efficiency of the operating system. We achieved 500 packets/s.)*

The interface computer. Could be merged with the node. Tentative and sparse information in this section.

The local network. Includes both interfaces for simple terminals and packet level interfaces for intelligent devices (then thought of as 'host computers' since PCs were not envisaged). Description of a scheme using cascaded eight way character multiplexers which we used in our local network. A bit strange but it worked at megabit rates.

Network economics. Tentative and unreliable estimate giving 0.1 pence per packet for use of the high level network.

Conclusions and recommendations. 'The possibility of a common-carrier communication network for digital data has been explored and a particular system design has been described in this paper.' Five features of the network are then listed. These are the well known characteristics of packet switched networks. Priority for building such a public data network is urged.