#148

Analysis and Design Issues Addressed at ICCC '78*

Leonard Kleinrock

Computer Science Department, University of California, Los Angeles, California 90024, USA

The ICCC '78 meeting in Japan was notable for the large number of papers involved with the analysis and design aspects of computer networks. In this review, I comment on a number of these papers and address my comments strictly to the printed material in the Proceedings rather than to the oral presentations. As mentioned in the three other review papers, the meeting was very well organized and great attention was paid to detail. However, the way in which papers were assigned to sessions was not as carefully done as it should have been. Particularly, session titles often had little to do with the papers presented in that session. As a result of this problem, my comments in this review paper will not be presented session by session but rather will be grouped by topics of interest to computer communications and each of these topics will contain papers from the collection of sessions where appropriate.

Each of the following sections is devoted to a specific topic and these sections (topics), are organized in a logical fashing, the first sections beginning with the innermost issues of network design, succeeding sections moving out from the interior of the network toward more global issues and the final sections referring to large scale issues and services. They are organized as follows:

(1) Centralized Network Design; (2) Backbone Network Design; (3) Routing; (4) Flow Control; (5) Protocols; (6) Multi-Access Communication Protocols; (7) Satellite Packet Switching; (8) Integrated Networks; (9) Value Added Networks; (10) Applications.

* This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract MDA 903-77-C-0272.

Computer Networks 4 (1980) 175–185

1. Centralized Network Design

In this section we discuss four papers:

2231 "A Heuristic Approach to the Optimization of Centralized Communication Networks," C.R. De Backer

2232 "A Network Optimization by Mixed Integer Programming," S. Matsui

2234 "An Integrated Approach to Optimally Locating Network Access Facilities," W. Chou, F. Ferrante, M. Balagangadhar and L. Gerke

2433 "An Algorithm for Designing Multidrop Teleprocessing Networks," T. Yosnimura

The problem of designing centralized networks (also known as multi-drop networks) remains unsolved. In the papers of this section we see various approaches to taming this problem, mainly in the form of providing heuristics (suboptimal solutions).

Paper 2232 presents a procedure for finding the true optimal solution of a centralized network using mixed integer programming techniques. However the

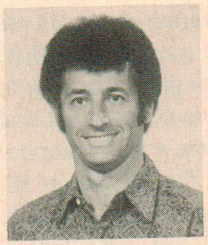

Leonard Kleinrock received the B.E.E. degree from City College of New York, NY, in 1957 and the M.S.E.E. and Ph.D.E.E. degrees from the Massachusetts Institute of Technology, Cambridge, MA, in 1959 and 1963 respectively. In 1963, he joined the faculty of the School of Engineering and Applied Science, University of California, Los Angeles, where he is now Professor of Computer Science. His research spans the fields of computer networks, computer systems modeling and analysis, queueing theory, and resource sharing and allocation in general. At UCLA, he directs a group in advanced teleprocessing systems and computer networks and is the Head of the Computer Science Department Research Laboratory. He is the author of three major books in the field of computer networks: *Communication Nets. Stochastic Message Flow and Delay* (New York: McGraw-Hill, 1964; also New York: Dover, 1972); *Queueing Systems, Vol. I: Theory* (New York, Wiley-Interscience, 1975); and *Queueing System, Vol. II: Computer Applications* (New York: Wiley-Interscience, 1976). He has published over 100 articles and contributed to several books. He serves as lecturer and consultant for many domestic and foreign corporations and governments and he is a referee for numerous scholarly publications and a book reviewer for several publishers. He is a Guggenheim Fellow, an IEEE Fellow, and has received various outstanding teacher and best paper awards, including the 1976 Lanchester prize for the outstanding paper in operations research, and the ICCC '78 prize-winning paper award.

problem is constrained to take a hierarchical form. In particular, most of the nodes are declared to be *slave* nodes and a few of the nodes are declared to be *master* nodes with one node declared to be the *center* of our centralized network. Slave nodes must be connected to master nodes and master nodes must be connected directly to the center, thereby preventing slave-to-slave connections and master-to-master connections. Of course this is a heavily constrained network design and therefore greatly reduces any computational effort to find that solution. The limited computation experience they quote seems to be reasonable for relatively small problems, for example, a problem involving only three master nodes. The paper would have been much stronger had the author given some indication for the way in which the computational effort grows with the number of master nodes and the number of slave nodes. This is, of course, the critical issue. Once one claims he has an optimal design procedure, the lack of such commentary is really quite unfortunate. (In the present paper a small sample of computational experience is provided).

Paper 2433 uses a heuristic approach for designing multi-drop networks. The approach is the usual one, namely, to select an initial network, go through local changes which create improvements, settling into a local optimal and then repeating the procedure until all the computational dollars are exhausted; at this point the best optimal solution is presented as the overall solution. The approach used by the author is to pick initial solutions in which terminals which are located close together initially are grouped together. The author is to be commended for giving the way in which the computation time for each local optimal grows, namely it grows as $n^{1.7}$ where n is the number of terminals; this is not especially fast growth at all, thereby indicating the effectiveness of his method at least in terms of computational cost. The author claims to give a 5–10% improvement over the Esau–Williams procedure; this is quite significant since experience shows that the Esau–Williams algorithm usually comes to within 5–10% of the optimum. Furthermore his procedure gets around the well-known degenerate cases in which the Esau–Willaims can be as bad as 50% worse than the optimum; this is rather interesting since one is concerned that one may be designing one of the degenerate networks and this author evidently gets around that problem. For some reason this paper was included in Session 243 whose title was "Fundamental Theories." It was not clear to this reviewer how one decides what is a fundamental theory and what is less than a fundamental theory in this field; it is simply clear that the papers in that session certainly do not seem more fundamental than many of the other papers at this conference.

Paper 2231 presents an approach which combines various earlier suboptimization procedures in an attempt to carry out an integrated design which includes the following design variables: topology, line capacities, number of concentrators, location of concentrators and line organization. These variables are then optimized so as to minimize cost while providing a reasonable response time. As the authors point out, the problem of optimizing one-level networks without concentrators has been solved in an exact way; but since the computational cost goes up quickly, that approach is useful only for small size problems. The authors of the current paper describe Kruskal's algorithm, Prim's algorithm, the Esau–Williams algorithm and the VAM algorithm. They then compare these various methods. Two-level designs are also discussed in this paper but the authors point out that they begin to fail when the problems get large, again due to excessive computing time (which unfortunately they do not present in a quantitative fashion). The authors of this paper do take advantage of the four earlier approaches to generate an initial solution which they modify by a parameterization process to achieve an improved solution. They claim that in approximately 85% of the cases they studied (30 small size problems with 5 and 10 nodes per problem) they obtained a solution within 2% of the computed optimum. They claim the computing times were reasonable but have failed to provide quantitative values for the computational time or even giving the way in which these computational times grow with network size.

Perhaps the most comprehensive paper in this session is paper 2234. These authors present an integrated approach to the optimal location of network access facilities (terminals, concentrators, etc.) and claim to apply this procedure both to the local access as well as the backbone network design procedure. However most of the paper is devoted to the local access problem. This paper is well written and presents an integrated approach. They include the usual techniques of adding, dropping, exchanging and merging links in the network design. They compare their approach to the MacGregor-Shen algorithm as well as the Chou-Kershenbaum unified algorithm; their current integrated approach seems to outperform both. What is even more lovely is that they give the total execution time for various-sized problems. They too find that decomposing the network into regions whereby near neighbors are connected to each other within the same region or nearby regions is an effective heuristic approach. Unfortunately they do not give an analytical expression for the way in which computation grows with the number of nodes, but at least they do give examples of a few network sizes.

In summary, we see that the war on local access network design goes on. We are slowly winning that war as we chip away at improved heuristics. Unfortunately the various approaches that we see still do not have a simple common means for comparison and this is especially frustrating for the reader. This reviewer recommends that any papers in the future which do not give a quantitative evaluation of the way in which their computational cost grows with the number of nodes not be accepted for publication.

2. Backbone Network Design

Here we discuss the design of backbone networks. The papers included in this section are as follows:

2233 "Topological Design of Computer-Communication Nets with Minimum Number of Through-Nodes," W.H. Kim and A.S. Donia

2333 "Extended Optimum Channel Capacity Assignment Problems for Message-Switching Networks," M. Komatsu, H. Nakanishi, H. Sanada and Y. Tezuka

2434 "Designing Reliable Packet-Switched Communication Networks," K. Maruyama

4122 "Reliability and Security: Primary Considerations in Network Design," R.D. Coop

4123 "Thirty Nine Steps to a Computer Network," J.L. Grange and M.I. Irland

4124 "Method of Synthesis of Communication Networks for Computers," V.G. Lazarev and N.Ya. Parshenkov

Paper 4124 addresses itself to network synthesis by presenting a heuristic algorithm for network design. This heuristic algorithm is stated but, unfortunately, the paper gives absolutely no computational experience with this procedure so we have no idea how effective it is or how computationally complex the algorithm is. As a result there is little more that one can say about the paper.

Paper 4122 addresses itself to reliability and security. It simply gives a framework for thinking about the issues of reliability and security, but presents no new insights and no new results. It is, however, an interesting introduction to the problem and presents an interesting list of considerations one must take into account when one approaches such issues.

Paper 2233 discusses network design using a graph theoretical approach to find networks with a minimum number of relay nodes. It is not clear why this particular problem is of great significance in the general scheme of things; one can see how this can be a consideration in their network design but it should not be the overriding concern. The algorithm presented is certainly elegant and interesting but suffers once again from the problem of giving no computational experience in terms of complexity.

Paper 2333 is a very nice paper which addresses the optimal channel capacity assignment in networks. This problem was first solved in [6] whereby the overall average network delay,

$$T = \sum \frac{\lambda_i}{\gamma} T_i$$

was minimized subject to a fixed cost constraint. Here λ_i = traffic on the ith channel, γ = total network traffic, and T_i = average time waiting for and using the ith channel as determined by queueing theoretic results. Seven years later in [15] the problem was generalized whereby the following function was minimized:

$$\left[\sum_i \frac{\lambda_i}{\gamma} T_i^\alpha\right]^{1/\alpha}$$

The authors refer to this as Problem A. Note that when $\alpha = 1$ the problem reduces to Kleinrock's original formulation and then the expression above is defined to be the overall network delay T. As is well known, the authors observe that Kleinrock's formulation is equivalent to minimizing the total average number of messages in the network since N_i, the average number of messages waiting for or using the ith channel, is given simply through Little's formula [8] by $\lambda_i T_i \triangleq N_i$. Of interest to the authors of this paper (seven years after [15]) is the minimization of the function.

$$\left[\sum_i N_i^\alpha\right]^{1/\alpha}$$

Again we see that when $\alpha = 1$ the problem reduces to Kleinrock's original problem. This new formulation is called Problem B. The first very nice result they give is to show

$$T^{*(\alpha)} = T^{*(1/\alpha)}$$

where $T^{*(\alpha)}$ is the resulting overall network delay when the capacity assignment solution to Problem B is plugged into the overall average network delay formula given above. Earlier Kleinrock had noted a similar relationship when $\alpha = 0$ and ∞ for Problem A [9]. A most interesting pair of results is found by the authors when the average length of all messages is the same:

$$C_i^{*(\alpha)} = C_i^{(1/\alpha)}\,, \qquad T^{*(\alpha)} = T^{(\alpha)}\,,$$

where starred quantities refer to the optimal solution of Problem B and the non-starred quantities refer to the optimal solution of Problem A, and further where C_i refers to the capacity of the ith channel. Note that whereas the capacity assignment in Problems A and B are quite different, the resulting average message delay is the same!

Paper 4123 is a delightful paper which presents some of the experiences of the CYCLADES project in France and the CCNG group in Canada, both of which cooperated in the development of the French CIGALE network. In particular, a simulation model was jointly developed and then was run by CCNG in an evaluation of the CIGALE network design. The simulator described is quite sophisticated and appears to have performed its main duty of identifying problems and successes with the proposed network design for CIGALE. The cooperation involved the investigation of the general network behavior and performance, transmission link protocols, adaptive routing and congestion (flow) control. Many of the observations made through their joint effort had earlier been made by the ARPANET researchers and the results given in this paper merely confirm what had already been known in many cases. For example, they describe simulation as well as analytic results which indicate that a limit should be placed on the length of each queue in the switch; this result had earlier been observed in [4] and [5]. The verification that the simple link protocol works had also earlier been observed. The fact that the overhead for packets was high, again had been observed in 1974 [7]. The fact that adaptive routing does in fact work was one of the key successes of the ARPANET and it is encouraging to see that their simulation also proved the stability of their adaptive routing procedure (they discuss an earlier fixed-routing algorithm whose entries were selected by an operator – they quickly found this presented problems, but of course this was well-known in the early days of the ARPANET). Their observation about setting thresholds for initiating routing updates was described many years ago in [2] and was named the Asynchronous Updating Algorithm; their approach is not unlike this algorithm. They refer to time-out periods and other numerical quantities throughout the paper but in many cases fail to give the value of these numerical quantities, an unfortunate and disturbing lack. An innovative flow-control algorithm which they carefully studied is the use of *choke* messages and various warning and alarm states. That study is extremely interesting and is supported in more detail by them in [14]. That study was extremely worthwhile. One further comment they make is how successful and useful it was to have two separate teams working on the same problem; they found that a close cooperation between two distinct teams generated useful criticisms and symbiotic progress. One real shortcoming of the paper is that they really have avoided a serious discussion regarding deadlocks; they introduce and study a variety of flow control and routing algorithms but fail to address the key problem of deadlocks, spending most of their discussion with regard to efficiency and degradations. Overall, however, this is a most enlightening paper and should be required reading for anyone seriously interested in network design.

Paper 2434 is concerned with the design of packet networks with a certain focus on the issue of reliability. This is a very good paper of the tutorial type with excellent pointers to the literature. This paper covers a broad area. It fails to go into enough detail on a variety of interesting design issues but nevertheless is worthwhile. They state the overall network design problem whereby they include delay not only for the channels but also the delay for nodal processing. Furthermore they address the reliability problem and point out that they have little choice but to use very approximate methods for computing node to node reliabilities. Their generalization of network design problems is to find the optimal topology, routing, link capacities, node capacities and priority discipline while minimizing network cost subject to constraints on throughput, network delay and *reliability*. The algorithm they introduce is called the TCPFA algorithm and, again, is not unlike previous algorithms. Many of the same observations were made in the paper by Gerla and Kleinrock in the second reference of paper 2434. The authors claim that "We have observed that the use of the tree network, which is designed subject to a set of design constraints, as the initial topology often reveals a less costly network than the one obtained from the use of a minimal spanning tree as the initial topology when the identical sequence of local optimization procedures is applied in both cases" – it would be very nice had the authors more clearly defined which was the initial tree topology that did so well.

3. Routing

There were very few papers devoted to routing procedures. The only one which we review here is the following:

2331 "Loop-Free Distributed Adaptive Routing," C. Neblock and U.W. Pooch

This paper, although included in a session entitled Analytic Techniques, is not an analytic paper at all. They claim to provide an algorithm which guarantees loop-free routing in distributed networks. Their simulation results are used to provide the performance of this algorithm. Their comments regarding the ARPANET for which they claim traffic is routed by a single path between delay estimate updates is somewhat misleading since multiple paths are permitted in the ARPANET due to traffic queued on the old paths from a previous update. Their whole approach to absolutely eliminating loops is also somewhat questionable since occasional loops are indeed needed when nodes and links fail. They discuss the advantages of BBN's old procedures but fail to discuss disadvantages which were mentioned [17]. Without giving proper reference to [2], the authors introduce what is equivalent to the shortest queue plus bias algorithm. They comment on Naylor's loop-free algorithm [16] and comment that the overhead is prohibitive without giving sufficient numerical data to support this. For their main examples they use the 19-node ARPANET described in [9] and reference their data to a certain percentage of the effective data rate used in that network; unfortunately when the effective data rate (RE) is equal to 100%, the load on the network is not maximum (i.e., the network did not saturate at RE = 1 in the original 19-node network whereas the reader of the current paper might easily be lead to believe that was the case). In fact they refer to 60% of network capacity when they really mean 60% of the effective data rate. Many of their constants were chosen as being equivalent to the ARPANET without their commenting on using the ARPANET values. They introduce the concept of a hierarchical routing procedure, an example of which had earlier been reported upon in [10]. The authors ought to be commended for presenting such a variety of simulation results which give one an excellent feeling as to how their algorithms perform. However, when all is said and done, the performance is not a significant improvement over existing algorithms and so one wonders how important the results really are.

4. Flow Control

In this section we present the papers devoted to flow control studies, all of which were in Session 142. The particular papers we review are

1421 "Analysis of Flow-Controls in Switched Data Network by a Unified Model," G. Pujolle

1422 "Comparison of Two Packet Network Protocols for Flow Control of Customer Virtual Circuits," F.R. Magee, Jr.

1423 "A Computer Network Flow Control Study," G.A. Deaton, Jr. and D.J. Franse

1424 "Evaluation of Flow Control Schemes for Packet Switched Network," I. Takenaka, K. Kawashima, M. Takeda and M. Nomura

We begin with paper 1423 in which the authors give a simulation study on certain aspects of network flow control. Basically they model parts of IBM's Systems Network Architecture (SNA) which they refer to as inbound pacing and which is simply a window type flow control between a terminal and a host across the network. Their major result is that network performance is improved with this form of flow control, which result offers no surprise. A serious problem with their paper is that it applies only to a single network of four nodes in a diamond shape topology. However, in defense of their work they do provide considerable detail in their simulation from an application to application point of view and thereby their results present a very realistic traffic environment. They show a variety of worthwhile curves supporting their results regarding the necessity of flow control. Furthermore they take the definition of power as presented by [3] and extend it to include buffer utilization and a measure of stability. Other extensions of notions of power have been seen in the literature, in particuler, [11,13];

Paper 1424 discusses four flow control schemes, namely

(1) Discard: Here the packet queue length is limited to a maximum value for each terminal and any packets arriving when the queue reaches this limit are discarded.

(2) WABT: This is the same as discard scheme except when a packet is dropped, a control message is

returned to the local source to prevent any further packet transmissions. When the queue length reduces below a given threshold, the source is then allowed to begin generating subsequent packets.

(3) GA: ("ga" stands for "go ahead.") The destination occasionally returns a "ga" packet (a token) which is really a bulk permit which permits the source to transmit k packets into the network. The destination issues these tokens based on its availability of free buffers.

(4) Permit: This basically is a simple permit scheme operating on a hop basis.

The authors have created a GPSS simulation program and have simulated a Japanese network. They claim that the permit scheme is deficient in that the throughput saturates very quickly as the traffic increases. Further they claim that "ga" works well only under a uniform traffic since with non-uniform traffic there is a hogging of buffers by certain flows; this result was earlier recognized in [4]. They further claim that there is a tradeoff between Discard and WABT depending upon the level of traffic but then finally conclude that WABT seem to be the most favorable. This paper is somewhat specialized and one wonders how general the results really are. These problems arise simply from having studied a single network under certain traffic assumptions.

Paper 1422 is not unlike the previous paper (1424) in which two particular protocols are now compared. The first is simply a window flow control scheme and the other is basically an on/off scheme as controlled by the destination which they name as the "restraint/resume" scheme. They find that the restraint/resume method uniformly gives lower overhead than the window method and is recommended in any network in which link-to-link rather than end-to-end packet acknowledgements are employed. Here too the paper suffers from the fact that the system for which they present their results is a simple tendem network of limited size and one wonders how general these results are in other network configurations and in cases where there is more than just tandem traffic competing with itself.

Paper 1421 is a very nice paper in which a unified model for flow control is studied. The author begins by showing that flow control is necessary (a fact we all know) and then identifies the various levels at which flow control can be exercised (node-to-node, host-to-host, and user-to-user). He then goes on to identify two handles often used in flow control schemes which, when adjusted, allow him to model many of the known flow control schemes in various existing networks. The first of these quantities is that of *thresholds* which operate on buffers, load and age of packets. The second refers to *limitations* either by *numbers* (global, on routes, on flows) or limitations on *time* (constant time, based on buffer state, based on line loading). He then goes on to show how the ARPANET, DATAPAC, CYCLADES, NPL and TRANSPAC all have flow control shcemes which can be modelled by a selection of various controls at various levels. The mathematical solution approach suggested by the author is a model using an equivalent station to represent a closed network system. This is a rather straightforward system whose behavior can be solved analytically. In exercising this model in various environments the author comes to the following conclusions: limitation in number by units of time should be used by small networks; credit schemes are useful for large networks; and preallocation is advisable when a known quantity of data must be transmitted. This is quite an interesting paper and is a useful analytical approach to the evaluation of flow control mechanisms.

5. Protocols

Here we have only one paper to report upon, namely,

4121 "Automatically Verified Data Transfer Protocols," J. Hajek

In this paper the author complains that current designers of computer communication protocols fail to have absorbed the lessons which were established in Dijkstra's paper from 1965 [1]. That earlier paper addressed itself to concurrent programming and is concerned with an absolutely safe solution of a mutual exclusion problem which certainly is the issue in verifiable data transfer protocols. The author of the current paper goes on to criticize a reliable transmission algorithm by Lunch and also criticizes NPL's "infallible" scheme. The author then talks about the "threshold of no return" and the "first threshold of no return" as important concepts in data transfer protocols. He shows that a self-synchronizing one-bit protocol can be designed, whereas it had been stated in previous papers that such a protocol was not possible.

6. Multi-Access Communication Protocols

In this section we include various access methods for distributed multi-access broadcast channels. We discuss the following papers:

3123 "Adaptive Retransmission Randomization Schemes for a Packet Switched Random Access Broadcast Channel," A. Fukuda, K. Mukumoto and T. Hasegawa

3221 "SRUC: A Technique for Packet Transmission on Multiple Access Channels," F. Borgonovo and L. Fratta

3222 "Simulation and Analysis of Satellite Packet Switching Computer Networks," W. Szpankowski, K. Ono and Y. Urano

3223 "Analysis of Framed ALOHA Channel in Satellite Packet Switching Networks," H. Okada, Y. Igarashi, Y. Nakanishi

It is interesting to note that all four papers were presented in sessions devoted to satellite packet switching (the word "packet" was spelled as "packed" in session 312!) and really these papers are not devoted to satellite systems alone but rather to the general multi-access schemes. Indeed there have been a number of papers in the literature devoted to this general problem of multi-access broadcast channels and the ones reported upon here represent a small portion of that literature.

The paper 3123 is a very nice paper related to the ALOHA access scheme in which an effective control policy is introduced and is referred to as an adaptive retransmission scheme. In particular, a packet which has met with a previous channel collision is retransmitted with probability p_k where p_k is a monotonically decreasing function of k. Their approach yields what appears to be a stable scheme and the issue of stability is clearly the key issue in ALOHA. They present an approximate analysis and a rather simplistic approach to the issue of stability although this system seems to have the capability for a very stable operation. They further give numerical results through simulation which tend to support many of their conclusions. They find that their approximation is accurate over a large range but becomes inaccurate near the point where the system goes unstable. They claim that slotted ALOHA is interesting because it has a relatively high capacity; one wonders if 37% maximum efficiency can honestly be called a high capacity. They do properly point out that one advantage of their adaptive ALOHA scheme is that it can be implemented *without* monitoring the channel history and without estimating the channel state. This is perhaps the most significant contribution of their paper and is worthwhile reading.

Paper 3321 presents a very nice adaptive access scheme which adjusts its parameters according to the measured load on the system and thereby behaves very much in the sense of the URN schemes [12]. The scheme here is called "split reservation upon collision (SRUC)". It is a combination of an ALOHA contention scheme with a superimposed reservation technique which operates automatically when collisions occur. The basic idea is to break a known population of n terminals into F groups and allow each group to compete only with other members in its own group. In those cases when collisions occur, the collisions are automatically resolved by a reservation procedure for members of the group that collided. There is overhead involved in this scheme due to the reservation and control information and it is this which is studied in the paper. This access scheme applies both to ground radio systems with very small propagation delays as well as to satellite systems with large propagation delays. It performs very nicely with regard to efficiency. Again the most interesting feature is that it is adaptive to load and therefore of significant interest. One problem is that one must know the number of terminals in the system and if this changes in a rapid fashion this could cause some synchronization and control problems.

Paper 3222 is of moderate interest. The authors introduce something known as periodic reservation ALOHA (PR-ALOHA) and give analytic results for this scheme. The analysis breaks down due to complexity and therefore simulation results are presented. Their techniques are very much similar to earlier techniques for analyzing ALOHA and slotted ALOHA systems. In the PR-ALOHA scheme the channel is divided into reservation slots and data slots. Reservation packets can collide in their scheme and only when a reservation packet is properly received will slots be put aside for that terminal's data packet transmission. This system is not unlike that of Robert's [18]. The authors find results for slotted and unslotted ALOHA which are the same as that found in earlier papers for these two systems. It is found that the PR-ALOHA system can achieve a higher troughput due to the reservation procedure but that for low traffic, additional delay is incurred over the usual schemes due to the reservation cycle; again, this is very much like the results obtained in [18].

Paper 3223 is a good paper in which a scheme called Framed ALOHA is studied. In this scheme a frame is defined and is at least equal in length to that of a round-trip propagation delay. The scheme is very much like slotted ALOHA except that a user is restricted from sending more than one packet in a given frame. The authors proceed to analyze the behavior of this system but they do so under the very simplified assumption of zero propagation delay; this makes the analysis quite manageable whereas the inclusion of a round-trip propagation delay renders this analysis (and that of many other satellite access schemes) quite impossible. They reach the following conclusions: They found through simulation that the propagation delay seemed not to impact the overall throughput characteristics (not surprising). Their main conclusion is that whereas the average throughput of Framed ALOHA is the same as that of slotted ALOHA, the variance of the throughput is reduced by a factor equal to the frame length as compared to slotted ALOHA – this is quite an impressive result. The authors give no results for delays with their system.

7. Satellite Packet Switching

In this section we study the various aspects of satellite packet switching and in particular discuss papers

3121 "International Broadcast Packet Satellite Services," E.V. Hoversten and H.L. Van Trees

3122 "On the Measured Performance of Packet Satellite Access Schemes," L. Kleinrock and M. Gerla

3224 "Evaluation and Measurement of the Performance Parameters for a High-Speed Computer-to-Computer Satellite Link," S.W. Ulmer

Paper 3121 is an extremely useful paper in which is given a description of a new international satellite service to provide international computer communications. It discusses the use of a *broadcast* packet satellite channel in order to provide computer communications between incompatible networks and/or terminals. The way in which a broadcast satellite can be used for multi-access communications is described and the general configuration is given. Details are given for the earth station equipment including such things as a satellite interface message processor (SIMP) and the satellite network interface processor (SNIP). The details of the terrestrial interface are also described in some detail and the flexibility in growth and nature of communication services is well described. This is a very excellent paper introducing the concept of a new satellite service which is currently under experimentation.

Paper 3122 is an excellent paper (if you are willing to believe my biased view). The paper presents some of the underlying tradeoffs in multi-access broadcast satellite channels and then proceeds to describe the ARPA supported satellite network experiment (SATNET). What is especially nice about this paper is that after the access protocols are described, analytical models are presented in order to predict performance in a way which assists the measurement experiments. The results of these experiments are presented in graphical form and illustrate a variety of throughput, delay and stability performance profiles. In particular it is shown that the control procedure used to stabilize S–ALOHA is quite effective and readily adapts to changing load conditions.

Paper 3224 describes a joint effort between IBM, COMSAT, and the French PTT in the use of the Symphonie satellite. The paper describes the performance parameters, the access control procedure and the measurement subsystem, but presents absolutely no measurements or data – an unfortunate omission and one which leaves the reader somewhat dissatisfied for lack of hard data; instead the reader is asked to find the detailed numerical results in a final report to the FCC!

8. Integrated Networks

In this section we discuss integrated networks, referring both to voice and data integration as well as to data communication media integration. We discuss papers

2432 "Queueing Analysis of Integrated Switched Networks," M. Watanabe, H. Miyahara and T. Hasegawa

3131 "Integrating Voice and Data Traffic in a Broadcast Network Using Random Access Scheme," A. Pan

3132 "Evaluation of Different Data Communication Media within an Integrated Telecommunication System," R. Parodi, E. Divano and L. Musumeci

Paper 2432 provides a queueing analysis of integrated switched networks where data packets mix with voice packets. The analysis has to do with dividding time in a dynamic fashion between these two traffic streams. Theirs is an approximate analysis which appears to be rather good. They make certain assumptions within the paper which still require proof (for example, the fact that the zeros of the denominator of Equation 2.10 are unique) but in general the analysis is complete. It would have been interesting had they compared their results to that of Gaver (see [9]) in which he studied remote terminal access similar to that which they are studying in this paper.

Paper 3131 also discusses an integrated voice and data network in a broadcast channel environment. They take advantage of the silence in voice in an adaptive fashion as one has seen in many of the other modern integrated systems. Among other things they prove the now well-known result that bigger systems are better in terms of resource sharing. They fail to refer to some of the ARPANET experiments in the behavior of voice packets in a packet network and this would have been a useful comparison had they carried it out.

Paper 3132 is a semi-interesting comparison of packet switching and circuit switching. They provide a rather simple framework for this comparison but when they are all done with numerous curves comparing the performance, one is left with an uneasy feeling that some of the basic cost assumptions have been unclearly stated. They do conclude that "There is no doubt that long distances favor packet switching." Further, they agree that interactive applications over relatively short distances are appropriate for packet switching.

9. Value Added Networks

In this section we describe the papers presented in session 321 entitled Packet Switching, but more appropriately should have been entitled Value Added Networks. We include the following papers:

3211 "DATAPAC Subscriber Service Update System Architecture and Operation," D.L. Jeanes, A.K. Trivedi, W.J. Older and K.G. Hayward

3212 "DDX Packet Switched Network and its Technology," M. Matsumoto, J. Mizusawa and H. Ohnishi

3213 "Development of a Packet Switching Node Prior to Opening of the French PTT Network, to Help Future users Debug their Network Access Software," C. Bertin, J.C. Adam, P. Guinaudeau and Y. Matras

3214 "Packet Switching as an Additional Public Service Offered by the Deutsche Bundespost," F. Hillebrand

Paper 3211 is devoted to a description of the service data update facility for the Datapac network. Indeed the idea is to provide sufficient subscriber service data to facilitate accounting in the network. The paper further describes the SL-10 node hardware, the subscriber services subsystem and node database, the support for the centralized database and the operation of the update system. This service database is a distributed database and allegedly provides improved performance and reliability. Indeed the network itself provides the facilities for distribution and updating of the service data. This is an interesting paper and worthwhile reading.

Paper 3212 describes specific features of the Japanese DDX Packet Switched Network, devoting itself to network access protocols, network configuration, hardware organization and network quality and maintenance facilities. Unfortunately the general tone of the paper would lead the uninitiated reader to believe that packet switching was invented in Japan! (Indeed, of the five references only one is a non-Japanese reference and it is to the CCITT.) Nevertheless, the paper is worthwhile reading and shows the way in which many of the currently operating networks had their influence on the design of DDX-1 and the DDX-2. Their use of the English language is rather amusing in places; for example, "Standby equipments in PS are confirmed in their sanities by . . . is checked by sending health-check packet by PS to . . . ." Among the operational facilities, they discuss congestion control and routing and come to the correct conclusion that a networkwide congestion control method is necessary; unfortunately the details of that control are vaguely stated in the paper (also their discussion of alternate routing used in DDX is quite vague).

Paper 3213 has one of the longest titles in the conference, but presents a worthwhile digest of information. In particular the French have developed an effective tool which allows users to experiment on what they refer to as a mini-TRANSPAC network to assist in the debugging of user application software prior to the actual connection to TRANSPAC. This is an example of a tool using measurements and hooks which in some sense has been neglected in some of the network implementations we have seen around the world. The authors refer to some successful experimentation with this tool.

Paper 3214 is a paper devoted to the packet switching service to be offered by the Deutsche Bundespost. At last Germany has come to realize the advantages of packet switching! Unfortunately the paper is written at a very high level and therefore lacks detail. One wonders now the tariffs and the charging structure compare to some of the existing public networks. The author points out that a large portion of the cost is in accessing the network and of course he is correct in this conclusion. He suggests that an intelligent microprocessor controlled concentrator be used to alleviate some of this cost.

10. Applications

In this section we discuss some applications to character recognition and filing systems. In particular we study the following papers:

3133 "A Character Recognition Communication System," H. Kaiyoh, K. Komori and K. Nakajima

3134 "The Speech Filing Migration System," D. Zeheb and S.J. Boies

2332 "Trade-off of File Directory Systems for Data Base," H. Tominaga, S. Tajima, K. Saito

Paper 3133 describes a character recognition communications system in which characters are recognized *remotely*. They provide some very nice experimental results in the use of the system and show some fairly efficient recognition capability for their proposed system.

Paper 3134 describes how digitally recorded speech may be stored in a hierarchically organized memory and retrieved. It is surprising that the authors did not refer to the enormous literature available on paging algorithms which seems to be exactly the kind of problem they are facing in their storage management. Nevertheless it is an interesting concept and makes good reading.

Paper 2332 presents a relatively uninteresting model describing the tradeoff among communication, computing, storage and total operating cost in a file directory system. Their assumption that the more frequently a file is queried, the more frequently a file will be updated is somewhat questionable and it would have been nice had they supported that assumption with some measurements. The cost model they set up is quite straightforward and the numerical results they show are keyed to a particular Japanese university network which they use as an example, thereby leaving the reader with a question as to how general the results are; in particular, specific constants are assumed with no real support for their values. Whereas this result may be useful for the particular network it was not clear that this was a universally general paper.

Conclusions

As we look back over the wealth of papers presented at the ICCC which were devoted to analytic aspects of computer network design, we find that there were really some very excellent papers with some worthwhile results. More than that, it is clear that there are a number of excellent groups working on these problems around the world and it is this which makes the future of network analysis, design and implementation so encouraging, we observe that many of the problems have been reasonably well-solved, so far as network design is concerned, but there still remains a large number of as yet unresolved problems to which we must address our energies in the near future. These include, among other things, flow control issues, congestion and routing control issues, protocol design, specification and verifiability of protocols, integrated network design and evaluation (both voice-data, and packet-circuit hybrid networks), process-to-process communication, distributed multi-access broadcast channel evaluation, local network architectures and their evaluation, distributed system control and evaluation, and overall network architectures. It is clear that the future of packet networks will depend not so much on the solution to these technical problems but rather on the solution of some of the softer problems which have been addressed at this conference as well, namely political, sociological, legal, tariff, regulatory, and standards issues. I leave these problems to my colleagues.

References

[1] E.W. Dijkstra, "Solution to a Problem in Concurrent Programming Control," Comm. ACM 8, 9, September 1965, p. 569.

[2] G. Fultz, and L. Kleinrock, "Adaptive Routing Techniques for Store-and-Forward Computer-Communication Networks," Conference Record, International Con-

ference on Communications, Montreal, Canada, June 1971, pp. 39-1 to 39-8.

[3] A. Giessler, J. Hanle, A. Konig and E. Pade, "Free Buffer Allocation – An Investigation by Simulation," Computer Networks, Vol. 1, No. 3, July 1978, pp. 191–204.

[4] F. Kamoun, "Design Considerations for Large Computer Communication Networks," Ph.D. Dissertation, Computer Science Department, University of California, Los Angeles, UCLA-ENG 7642, April 1976.

[5] P. Kermani, "Switching and Flow Control Techniques in Computer Communication Networks," Ph.D. Dissertation, Computer Science Department, University of California, Los Angeles, UCLA-ENG 7802, February 1978.

[6] L. Kleinrock, Communication Nets, Stochastic Message Flow and Delay, McGraw-Hill, New York, 1964. Out of print. Reprinted by Dover Publications, 1972.

[7] L. Kleinrock and W. Naylor, "On Measured Behavior of the ARPA Network," AFIPS Conference Proceedings, Vol. 43, National Computer Conference, Chicago, Illinois, May 1974, pp. 767–780.

[8] L. Kleinrock, Queueing Systems, Vol. I: Theory, Wiley-Interscience, New York, 1975.

[9] L. Kleinrock, Queueing Systems, Vol. II: Computer Applications, Wiley-Interscience, New York, 1976.

[10] L. Kleinrock and F. Kamoun, "Hierarchical Routing for Large Networks, Performance Evaluation and Optimization," Computer Networks, Vol. 1, No. 3, January 1977, pp. 155–174.

[11] L. Kleinrock, "On Flow Control in Computer Networks," Conference Record, International Conference on Communications, Vol. 2, Toronto, Danada, June 1978, pp. 27.2.1–27.2.5.

[12] L. Kleinrock and Y. Yemini, "An Optimal Adaptive Scheme for Multiple Access Broadcast Communication," Conference Record, International Conference on Communications, Vol. I, Toronto, Ontario, June 1978, pp. 7.2.1 to 7.2.5.

[13] L. Kleinrock, "Power and Deterministic Rules of Thumb for Probabilistic Problems in Computer Communications," Conference Record, International Conference on Communications, Boston, Massachusetts, June 1979, pp. 431.1 to 431.10.

[14] J. Majithia, M. Irland, J. Grange, N. Cohen and C. O'Donnell, "Experiments in Congestion Control Techniques," in: Flow Control in Computer Networks, J. Grange and M. Gien, eds., IFIP, North-Holland Publishing Company, Amsterdam, 1979.

[15] B. Meister, H. Muller and H. Rudin, "New Optimization Criteria for Message-Switching Networks," IEEE Transactions on Communication Technology, Vol. COM-19, June 1971, pp. 256–260.

[16] W. Naylor, "A Loop-Free Adaptive Routing Algorithm for Packet Switched Networks," Proceedings of the Fourth Data Communications Symposium, Quebec City, Canada, October 7-9, 1975, pp. 7.9–14.

[17] H. Opderbeck and L. Kleinrock, "The Influence of Control Procedures on the Performance of Packet-Switched Networks," National Telecommunications Conference Record, San Diego, California, December 1974, pp. 810–817.

[18] L.G. Roberts, "Dynamic Allocation of Satellite Capacity through Packet Reservation," AFIPS Conference Proceedings, Vol. 43, National Computer Conference, New York, 1973, pp. 711–716.