Hierarchical Use of Dedicated Channels *

Gideon Y. Akavia **

Center for Military Analyses (CEMA), P.O. Box 2250, Haifa, Israel

Leonard Kleinrock

Computer Science Department, University of California, Los Angeles, CA 90024-1596, U.S.A.

Received 12 November 1986
Revised 1 March 1988

We consider efficient organizations for communication resources which are accessed by a large number of geographically distributed terminals. Developing a model for systems built with dedicated line channels, we answer the following questions: What is the role of hierarchies in organizing large communication nets? How should a large network be decomposed into smaller parts? What cost vs. performance gains can be achieved by such a decomposition?

Assuming that the traffic to be carried and the mean response time are specified and that the goal is to minimize the necessary cost, we define *burstiness* and find the following: Dedicating channels is reasonable when the traffic is steady (i.e., not bursty) but when the traffic is bursty, the cost of simple dedicated-channel systems grows too fast with the number of terminals. By introducing *regular* hierarchical structures we show that the cost of bursty systems can significantly be reduced. The *optimal* structure must be balanced. The optimal number of levels and the ratio of the contribution of the different levels to both cost and delay is simply determined by a few key systems parameters.

Keywords: Hierarchical Network, Dedicated Channel, Large Communication Network, Burstiness, Optimal Network, Regular Hierarchical Structure.

* This research was supported by the Defense Advanced Research Projects Agency of the Department of Defense under Contract No. MDA 903-87-C-0663.

** Presently visiting the Center for International Security and Arms Control (CISAC), Stanford University, 320 Galvez Street, Stanford, CA 94305, U.S.A.

North-Holland
Performance Evaluation 9 (1988/89) 135–142

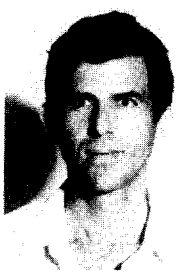

Gideon Y. Akavia has a Ph.D. in Computer Science from the University of California in Los Angeles. He was with the Faculty of the Israel Institute of Technology, and since 1980 has been with the Center for Military Analyses (CEMA), Haifa, Israel. He is presently visiting Stanford University, as a Science Fellow in the Center for International Security and Arms Control. His research interests focus on military applications of advanced technology.

Leonard Kleinrock is a professor of Computer Science at the University of California, Los Angeles. He received his B.S. degree in Electrical Engineering from the City College of New York in 1957 (evening session) and his M.S.E.E. and Ph.D.E.E. degrees from the Massachusetts Institute of Technology in 1959 and 1963, respectively. While at MIT, he worked at the Research Laboratory for Electronics, as well as with the computer research group of Lincoln Laboratory in advanced technology. He joined the faculty at U.C.L.A. in 1963. His research interests focus on local area networks, computer networks, performance evaluation and distributed systems.

He has had over 150 papers published and is the author of five books—*Communication Nets: Stochastic Message Flow and Delay* (1964), *Queueing Systems, Vol. I: Theory* (1975), *Queueing Systems, Vol. II: Computer Applications* (1976), *Solutions Manual for Queueing Systems, Vol. I* (1982), and, most recently, *Solutions Manual for Queueing Systems, Vol. II* (1986). Dr. Kleinrock is co-director of the U.C.L.A. Computer Science Department Center for Experimental Computer Science and is a well-known lecturer in the computer industry. He is the principal investigator for the DARPA Laboratory contract at U.C.L.A.

Dr. Kleinrock is a member of the National Academy of Engineering, a Guggenheim Fellow, an IEEE Fellow, and a member of the IBM Science Advisory Committee, and in 1986 he became a member of the Computer Science and Technology Board of the National Research Council. He has received numerous best paper and teaching awards, including the ICC 1978 Price Winning Paper Award, the 1976 Lanchester Prize for outstanding work in Operations Research, and the Communications Society 1975 Leonard G. Abraham Prize Paper Award. In 1982, as well as having been selected to receive the C.C.N.Y. Townsend Harris Medal, he was co-winner of the L.M. Ericsson Prize, presented by His Majesty King Carl Gustaf of Sweden, for his outstanding contribution in packet switching technology. In July 1986, Dr. Kleinrock received the 12th Marconi International Fellowship Award, presented by His Royal Highness Prince Albert, brother of King Baudoin of Belgium, for his pioneering work in the field of computer networks.

Dr. Kleinrock is also founder and chairman of Technology Transfer Institute, a computer/communications seminar and consulting organization located in Santa Monica, California.

1. Introduction

Designing a communication network for a given traffic requirement consists of balancing cost and performance. Faced with the task of analyzing networks we must abstract the relevant features of traffic, performance and cost in order to arrive at a manageable model. In this paper we develop such a model and use it to answer the following questions: What is the role of hierarchies in organizing large communication nets? How should a large network be decomposed into smaller parts? What cost vs. performance gains can be achieved by such a decomposition?

In order to be able to treat the issue of hierarchies in large networks, we must drastically simplify all other aspects of our model. To characterize traffic we shall assume that messages originate at m different sources (buffered terminals), that the appearance of messages at each source is an independent Poisson process with rate S/m messages per second, and that the length of messages has an exponential distribution. Let us choose the information unit so that the average length of a message is equal to 1; this is simply a convenient normalization, which is equivalent to measuring communication capacity in messages (of an average length) per second, instead of measuring in bits per second. Assume all messages are directed to one destination (computer), which we shall sometimes call the station. We shall characterize performance by T, the average total time spent by a message in the system.

If all terminals are in one place, the best thing is to connect them all to the destination by a single communication channel. Each message will join a queue at the terminals' end of the channel, and when its turn comes, will be transmitted to its destination. We thus have a classical M/M/1 queueing system [6] whose service rate is C, i.e., the capacity (bandwidth) of the common communication channel, measured in messages per second. The necessary C in such an M/M/1 system when S and T are given is

$$C_{\mathrm{M/M/1}} = S + \frac{1}{T}. \tag{1}$$

If terminals are not co-located and cannot simply join a queue and share a single channel, we may dedicate a subchannel with capacity C/m to each. Each message will therefore have to pass through one of m identical queueing systems (with arrival rate S/m and service rate C/m each). The total channel capacity needed in such a network with m dedicated subchannels is

$$C_{\mathrm{ded}} = S + \frac{m}{T}. \tag{2}$$

Dividing equations (1) and (2) we get

$$\frac{C_{\mathrm{ded}}}{C_{\mathrm{M/M/1}}} = \frac{ST + m}{ST + 1}. \tag{3}$$

We shall call the inverse of ST the 'burstiness' [2] of the system. When ST is small ($ST \ll 1$), the system is bursty. When ST is large ($ST \gg 1$), the system is steady. Using these terms we see, from equation (3), that when the system is very bursty, using m dedicated channels requires m times the amount of total channel capacity as compared with sharing a single M/M/1 channel. When the system is very steady, m dedicated channels are almost as good as a single M/M/1 channel. Definitions equivalent to our burstiness were introduced independently by others [4,9]. This is not surprising, since ST is the only dimensionless number one can form with S and T.

When the traffic is bursty, there are only a few messages in the system. There is little congestion, and the delay suffered by messages is mainly determined by the time necessary to transmit them. The communication resource is only lightly utilized in a bursty system. When the traffic is steady, the communication resource is heavily utilized and the delay is mainly determined by the congestion.

It should be noted that we assume arrivals of messages are independent, and use the term 'bursty' to characterize ST. Others sometimes use the term 'bursty' to describe a non-Markovian system, in which separate messages tend to congregate and arrive together.

When terminals are distributed in space, it is very hard for them to perfectly share common communication resources. People have therefore developed many schemes, often called access modes, for deciding which terminal will use which part of the communication resources at a given time (for examples, see [1,3,7,8,10,12]).

The performance of many of the access modes is extremely hard to obtain in an analytic way because they involve complex systems of interacting queues. While it is often easy to evaluate an access mode for a small range of parameters by simulation, it is hard to use simulation to gain insight into which access mode is best for which range of parameters. We shall treat analytically very large and very bursty communication networks using dedicated line channels, and shall answer the following question: For a given traffic and required performance, can the cost of a very bursty system be reduced by a hierarchical organization?

We shallthat assume that m, the number of terminals, is very large, that terminals are uniformly distributed in their very large geographic region, and that all terminals contribute equally to the traffic. The reason is that we are interested in hierarchies that arise in the design process, and not in hierarchies that are imposed by the geography of the region and by uneven traffic requirements. It is also often true that the uniform case is the worst case for a distributed system: If traffic is especially concentrated in some terminals or regions, then the system would be less distributed.

We shall assume that the cost of a line channel of length l and capacity C is $l^a C^b$. When there is an economy of scale in buying capacity, we have $b < 1$. When there is no economy of scale, we have $b = 1$. There is no reason to consider diseconomy of scale, and there is no reason, therefore, to consider the case when $b > 1$. When the cost depends linearly on length, we have $a = 1$. When the cost of equipment at the ends of the lines is significant, it is natural to assume that $a < 1$. We have adopted the simplest possible assumptions on cost that still represent economy of scale and line vs. equipment cost. We cannot generalize our results to other, more complicated, cost functions.

Real systems are built slowly. Investments have to be based on estimates of future demand, and the demand in the future is influenced by the existence of the system and the quality of service. We shall ignore this interaction over time, and assume that our systems are built in order to satisfy the known demand and service requirements at a given time.

In Section 2 we solve the problem of allocating resources to subsystems, a problem that must be faced whenever a system is decomposed. In Section 3 we propose and analyze a heuristic for decomposing, a heuristic that leads to regular hierarchical structures.

2. Decomposition and resource allocation

Assume we have a procedure for designing a very bursty line-based communication system, given m, the number of equally talkative and uniformly distributed terminals. Denote by L the size of the region over which terminals are distributed. The length of all lines connecting terminals to the one station is proportional to L, and the cost of these lines is therefore proportional to L^a. When the traffic is very bursty, the typical capacity of every line channel must be proportional to $1/T$, and it follows that the total system cost is proportional to $1/T^b$. The total cost D can therefore be written, without loss of generality, as

$$D = \frac{L^a}{T^b} f. \tag{4}$$

Given our assumption on the cost of individual lines, the dependence of D on L and on T is an inevitable result of the traffic requirements, i.e., of wanting to communicate across distances (that are typically L) over lines (whose capacity must typically be $1/T$). The f appearing in equation (4) shows how the system cost depends on its being distributed, and characterizes the quality of the design procedure. f contains some geometric constants, and a dependence on the number of terminals m.

For example, consider star networks in which every terminal is connected by an individual line directly to the station. The cost of such networks is proportional to the number of lines, i.e., to the number of terminals. Assume that terminals are uniformly distributed in a disk of radius L, with the station in its middle; assume all terminals are equally talkative, and that $a = b = 1$. If we assign to all lines capacity $C = 1/T$ and make the average terminal-station delay the same for all terminals, the cost will be

$$D = \frac{L}{T}\frac{2}{3}m.$$

The geometric constant $\frac{2}{3}$ is the average distance from a unit disk to its center. If we allocate capacity to minimize the cost and demand only that the delay over all terminals will be T, the cost will be slightly smaller:

$$D = \frac{L}{T}\frac{16}{25}m.$$

The geometric constant here is $\frac{16}{25}$, the square of the average square root of the distance from a unit disk to its center. (See the following discussion of optimal allocation and equation (6)).

We shall usually ignore the geometric constants, and address the dependence on m. How fast does f grow with m? Must it grow that fast? Can we reduce the cost of a communication system by decomposing it into small parts, and by applying the given design procedure to each part separately? How should we decompose a large system and how should we allocate resources to the different subsystems? We shall start with the latter question. Assume the cost of the jth subsystem is given by equation (4), i.e.,

$$D_j = \frac{L_j^a}{T_j^b} f_j$$

and that the total system cost is $D = \sum_j D_j$. Assume that the delay measure T is given by the following weighted average

$$T = \sum_j S_j T_j / S, \tag{5}$$

where S_j is the traffic carried by the jth subsystem and S is the total traffic. The problem of allocating resources to subsystems can be posed as an optimization problem: choose the T_j (and the D_j) in order to minimize D. But, the T_j cannot be chosen freely. Since T is given, their choice is constrained by equation (5). This problem of constrained optimization can be solved very simply, using a Lagrangean multiplier. When the optimal choice of T_j is made, the total cost is given by

$$D = \frac{1}{(ST)^b} B^{b+1}, \tag{6}$$

where

$$B = \sum_j \left(L_j^a S_j^b f_j \right)^{1/(b+1)}.$$

Minimizing the cost of a hierarchical structure often involves minimizing B given in equation (6), which we shall call the B-term.

When resources are allocated to subsystems in the optimal way, which leads to (6), we also get

$$\frac{D_j}{D_k} = \frac{S_j T_j}{S_k T_k} = \left[\frac{L_j^a S_j^b f_j}{L_k^a S_k^b f_k} \right]^{1/(b+1)}. \tag{7}$$

That is, the contributions of subsystems to the delay measure and to the cost are directly proportional to their contributions to the B-term. (Equations (6) and (7) were obtained while minimizing D given T, but they are also valid when minimizing T given D.)

When our subsystems consist of a single line each, equation (6) is very similar to Kleinrock's optimal capacity assignment [5], with the following difference: By restricting ourselves to very bursty traffic, we can handle cost functions with any b, not just the $b = 1$ case. While only the average delay T appears in our formulas, our main results will remain valid when the variance, range or distribution of acceptable delay values is specified in addition to the average delay. For example, Meister et al. [11] propose a performance measure that can influence the variance of delay. If their $T^{(k)}$ is used as a performance measure and b/k is substituted for b, equations (6) and (7) remain valid with only slight modifications.

3. Regular hierarchical structures

Having decomposed a communication system, equation (6) gives a way to allocate resources to its various parts. We do not know which is the optimal way to decompose a large system for our goal of minimizing cost, so we shall use a heuristic. To introduce it, consider the following two-level structure: Assume that the m terminals are uniformly distributed in a region of the plane, and divide this region into P regions. Place a concentrator in the middle of each region. Connect all P concentrators to the station according to a given design procedure, and connect all terminals in a given subregion to 'their' concentrator according to the same design procedure. For simplicity of our formulas we shall assume that all subregions have the same shape as the original region, and ignore the geometric constraints that depend on this common shape.

We shall call this hierarchical system a two-level regular hierarchical system, where the word 'regular' refers to the fact that all regions are of the same size and shape, and that all concentrators are placed in the middle of their regions. We shall call the communication subsystem connecting concentrators to the station the top level, and the subsystem connecting terminals to concentrators the bottom level. The top level consists of a network within the P concentrators acting as terminals, and the bottom level consists of P networks with m/P terminals each.

Let L be the typical length of the original region. The typical length of each one of the P subregions is $L/\sqrt{P}$, and the total traffic arriving at each concentrator is S/P. Applying (6) to both levels we find that the contribution of the bottom level to the B-term is

$$P\left[L^a(1/P)^{a/2}(S/P)^b f(m/P)\right]^{1/(b+1)},$$

where we have explicitly shown the dependence of f on m/P, the number of terminals in every subregion. The contribution of the top level to the B-term is

$$\left[L^a S^b f(P)\right]^{1/(b+1)}.$$

Summing these two gives the B-term of the two-level regular hierarchical system:

$$B=\left[L^a S^b\right]^{1/(b+1)}\left[f(P)^{1/(b+1)}+P^{(1-a/2)/(b+1)}\right.$$
$$\left. f(m/P)^{1/(b+1)}\right]. \qquad (8)$$

Which P will give the least cost two-level system? To find the best P that will minimize B we must say something about the f-function. How fast does the cost grow when m grows? Does $f(m)$ grow like an exponent of m, like a polynomial in m, or like the logarithm of m? The cost of the simplest possible star network is proportional to m, so there is no need to consider exponential growth. If cost grows like the logarithm of m, we are in good shape and cannot improve matters by regular hierarchical structures. When $f(m)$ grows like a polynomial in m, it is dominated, for large m, by its leading term. We shall, therefore, only treat the case when, for large m, the following is a good approximation:

$$f(m)=m^g. \qquad (9)$$

If g is small, our original design procedure is already quite good. We shall show below that if g is larger than $1-a/2$, then a two-level hierarchical system with the best P will be better than a one-level system. We shall now calculate this best P.

Assuming that P satisfies $m \gg P \gg 1$, so that both P and m/P are large compared to unity, we can substitute (9) into (8) and get

$$B=\left[L^a S^b\right]^{1/(b+1)}$$
$$\left[P^{g/(b+1)}+P^{(1-a/2)/(b+1)}(m/P)^{g/(b+1)}\right]. \qquad (10)$$

Differentiating B with respect to P we see that $\mathrm{d}B/\mathrm{d}P=0$ when

$$g^{b+1}P^g=\left[g-1+a/2\right]^{b+1}(m/P)^g P^{1-a/2}. \qquad (11)$$

Substituting P determined by (11) into (10) we see that the cost of the two-level structure, optimized with respect to P, is proportional to m^h, where

$$h=\frac{g^2}{2g-1+a/2}. \qquad (12)$$

When $g>1-a/2$, we have $g>h$. That is, when using the best P, as given by (11), we have a two-level structure whose cost grows with m more slowly than the cost of the one-level structure. When $g>1-a/2$ and $m \gg 1$, our use of approximation (9) is consistent, since our best P does satisfy $P \gg 1$ and $m/P \gg 1$. We can summarize the above discussion of two-level regular hierarchical systems by the following result.

3.1. Theorem. *A design procedure whose cost is proportional to m^g where $g>1-a/2$ can be improved for large m by applying it separately to each level of a two-level regular structure. The best P (number of groups) is given by* (11). *The cost of the resulting two-level structure is proportional to m^h, where h is given by* (12). *When the best P is used, the contribution of the two levels to the delay, to the cost and to the B-term satisfy*

$$\frac{T_{\text{top}}}{T_{\text{bottom}}}=\frac{D_{\text{top}}}{D_{\text{bottom}}}=\frac{B_{\text{top}}}{B_{\text{bottom}}}=\frac{g-1+a/2}{g}. \qquad (13)$$

Proof. Substituting (11) in (10) we get $B_{\text{top}}/B_{\text{bottom}}=(g-1+a/2)/g$. The other two equalities are true whenever capacity is optimally allocated, as shown in (7).

We shall paraphrase (13) by saying that the optimal two-level regular structure is *balanced*. The contribution of both levels to the delay and their share of the budget must be in the proportion given by (13). The right-hand side of (13) decreases when g decreases. We may say that when g is small, most of the system migrates to the bottom level, and that when g is small enough, two levels become unnecessary. P also decreases with g, and there will be less groups in the top level. This is satisfying intuitively, but the exact value of P should not be taken too seriously when small, since it was obtained assuming $m \gg P \gg 1$.

3.2. Example. When the original design procedure consists of building a star network, we have $g = 1$, and equation (12) reduces to $h = 2/(a+2)$. That is, the cost of the optimal regular two-level star system is proportional to $m^{2/(a+2)}$, while the cost of a one-level system is proportional to m. When $g = 1$, equation (13) reduces to

$$\frac{T_{\text{top}}}{T_{\text{bottom}}} = \frac{D_{\text{top}}}{D_{\text{bottom}}} = a/2.$$

If two levels are better than one, will more levels be even better? Equation (12) already contains the answer: Decomposing a given system into two levels and applying the original design procedure to each can be considered as a new design procedure. Applying this new procedure to two levels is equivalent to applying the original procedure to four levels. When $g > 1 - a/2$, it follows from (12) that $h > 1 - a/2$ and therefore four levels will be better than two when m is large enough. In general, let g_j be the power of m characterizing the resulting cost and f-function when the given design procedure is applied to 2^j levels. Equation (12) can be written as

$$g_{j+1} = \frac{g_j^2}{2g_j - 1 + a/2},$$

where g_0 is the power of m characterizing the direct application of the given design procedure to one level. It is easy to see that when $g > 1 - a/2$, the sequence $\{g_j\}$ is monotonically decreasing and converges to $1 - a/2$.

The argument of the previous paragraph has the flavor of an existence proof. It shows that by having enough levels the cost can be made to grow as an exponent of m arbitrarily close to $1 - a/2$. As m becomes larger, using more and more levels is justified. What is the best number of levels for a large but fixed m? To answer this question we must consider the constant coefficient multiplying m^{g_j}. This constant, which was ignored until now, grows with the number of levels, and therefore tempers the trend towards more and more levels.

The f-function and the cost of a system consisting of r levels, each of which is built according to a given design procedure, can be calculated explicitly. Let P_j be the number of terminals per group in the jth level, starting from the top. If an r-level system is optimal, then every two consecutive levels must be optimal as a group of two-level systems. Equation (11) can therefore be rewritten, when $g > 1 - a/2$, as

$$g^{b+1} P_j^{g-1+a/2} = [g - 1 + a/2]^{b+1} P_{j+1} g, \tag{14}$$

and equation (13) can be generalized into

$$\frac{B_j}{B_{j+1}} = 1 - (1 - a/2)/g, \tag{15}$$

where B_j is the contribution of the jth level to the B-term. From equations (14) and (15) we get the following result.

3.3. Theorem. *A design procedure whose cost is proportional to m^g where $g > 1 - a/2$ can be improved for large m by a multi-level regular organization.*

r, the best number of levels, is given by

$$r\frac{(b+1)}{g} \ln \frac{g}{g - 1 + a/2} = (1 - a/2) \ln m \tag{16}$$

and the cost of the system, when using this r, is proportional to

$$[m^{(1-a/2)/(b+1)} - 1]^{b+1}. \tag{17}$$

When the optimal number of levels is used, the number of lines in all groups at all levels is the same, and must therefore be given by $m^{1/r}$.

The proof is given in Appendix A.

According to our assumptions we have $a \leqslant 1$ and $b \leqslant 1$. When a is smaller, the best regular hierarchical system has fewer levels, since it is harder to save by shortening individual lines. When b is smaller, the best system has more levels and leads to larger improvements, since common large capacity lines become more economical.

3.4. Example. Let the given design procedure be to build a star network (that is, $g = 1$). Let a and b be equal to 1. From equation (16) we see that the optimal number of levels is given in this case by $\log_{16} m$, and that we should have 16 lines in every group. When using the optimal number of levels,

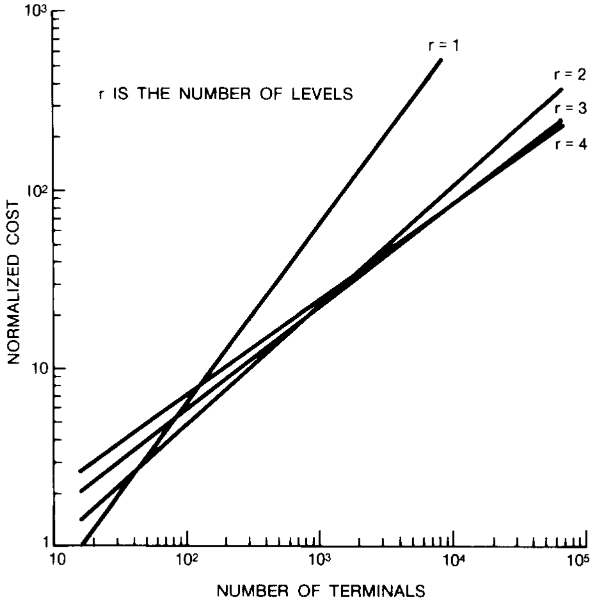

Fig. 1. Normalized cost when $a = b = g = 1$.

the cost of the resulting system is proportional to $(L/T)[m^{1/4} - 1]^2$.

The dependence of the B-term on r, the number of levels, is given by equation (A.6) (see Appendix A). We shall use the square of the expression in equation (A.6) as a normalized cost of various systems. Fig. 1 gives this normalized cost as a function of m, the number of terminals, when r, the number of levels used, is equal to 1, 2, 3, and 4. The normalized cost of the system with the optimal number of levels is $(m^{1/4} - 1)^2$ and, were it drawn on Fig. 1, it would have been very close to the lower envelope of the four straight lines drawn.

4. Conclusions

We have assumed that the traffic level and the necessary performance are specified, and that the goal is to fulfill these requirements with the least cost. Burstiness is defined and serves as a natural dimensionless number to characterize the requirements. We also assume that space is homogeneous: Terminal density and traffic requirements are the same everywhere. The validity of our results in the case of irregularity either in spatial distribution or in traffic requirements was not investigated. The cost of communication resources was modeled by simple power laws.

When the traffic is steady, the cost of simple one-level dedicated-channel systems is reasonable, since all channels will be well utilized. When you demand high performance, i.e., when you specify S and T such that $ST \ll 1$ and the system is very bursty, no congestion can be tolerated. Of course, such bursty systems are expensive, but sharing can reduce cost. Cost is reduced by sharing lightly utilized channels, and when the system is *very* bursty, sharing does not lead to congestion. Such sharing can reduce cost even if the technology has no inherent economies of scale.

To make sharing of dedicated channels possible, we introduce regular hierarchical structures. Our regular structures are obtained by dividing the terminal population into equal groups, and placing a concentrator in the center of each. Regular multi-level hierarchical structures can significantly improve the performance of bursty systems. The optimal structure is characterized by a balance principle that gives the ratio of investment in any two consecutive levels. Another characteristic of the optimal regular hierarchical structures is that channels are organized in small groups of equal sizes.

Our ability to analytically treat the structure of very large communication systems rests on two drastic simplifying assumptions: That traffic is very bursty, and that the dependence of cost on the number of terminals can be approximated by a simple power law ($f(m) = m^g$). Even if these assumptions are too drastic for some applications, we expect the main results, summarized in the last two sentences of the previous paragraph, to be of value as starting points for heuristics.

Acknowledgment

The contribution of the anonymous referees is greatly appreciated.

Appendix A

To simplify our formulas here let us rewrite equations (14) and (15) respectively as

$$tP_i^s = P_{i+1} \tag{A.1}$$

and

$$B_i/B_{i+1} = 1/x, \tag{A.2}$$

where

$$x = \frac{g}{g - 1 + a/2},$$

$$s = 1/x,$$

$$t = x^{(b+1)/g}.$$

Using $\Pi_i P_i = m$ and $\Sigma_i B_i = B$, we can solve equations (A.1) and (A.2) for P_i and B_i in terms of B, r, t, s, x, and m.

When $a \neq 2$, we get

$$P_i = t^{1/(1-s)}\left[m^{1-s}t^{-r}\right]^{s^{i-1}/(1-s^r)}, \tag{A.3}$$

$$B_i = x^{i-1}\frac{1-x}{1-x^r}B. \tag{A.4}$$

Ignoring geometric constants, we also know that the following must be true:

$$B_1 \cong \left[L^a S^b P_1^g\right]^{1/(b+1)}. \tag{A.5}$$

Using equations (A.3) and (A.4) in (A.5) we can get B as a function of m, r, and the constants t, s, and x. Isolating the dependence on r we get that B is proportional to

$$(x^r - 1)\left[m^{1-s}t^{-r}\right]^{(1/1-s^r)g/(b+1)}. \tag{A.6}$$

Differentiating, we find that B is minimized as a function of r, when

$$m^{1-s} = x^{r(b+1)/g}. \tag{A.7}$$

Substituting equation (A.7) in (A.6) we get that B is proportional to $x^r - 1$ and is therefore proportional to $m^{(1-s)g/(b+1)} - 1$. Since the cost is proportional to B^{b+1}, it follows that when the best r is used, the system cost is proportional to

$$\left[m^{(1-a/2)/(b+1)} - 1\right]^{b+1}. \tag{A.8}$$

Substituting equation (A.7) in (A.3), we also see that when the best r is used, the P_i do not depend on i, and they must therefore satisfy $P_i = m^{1/r}$.

While the best number of levels will depend on g, i.e., on the quality of the design procedure applied to each level, (A.8) shows that the system cost, when the best number of levels is used, is independent of g. For larger m we can also approximate (A.8) by $m^{1-a/2}$ and see that the growth with m of the best regular hierarchical system only depends on the length dependence of line cost, and hardly depends on the capacity dependence of line cost.

References

[1] N. Abramson, Packet switching with satellites, *AFIPS Conf. Proc., 1973 National Computer Conference, Vol. 42* (1973) 695–702.

[2] G.Y. Akavia, *Hierarchical Organization of Distributed Packet-Switching Communication Systems*, Ph.D. Dissertation, Computer Science Dept., Univ. of California, Los Angeles, CA, 1978.

[3] R. Binder, A dynamic packet-switching system for satellite broadcast channels, *Proc. IEEE Internat. Conf. on Communications*, San Francisco, CA (June 16–18, 1975) 41-1–41-5.

[4] M.J. Ferguson, A bound and approximation of delay distribution for fixed-length packets in an unslotted ALOHA channel and a comparison with time division multiplexing (TDM), *IEEE Trans. Commun.* **COM-25** (1977) 136–139.

[5] L. Kleinrock, *Communication Nets: Stochastic Message Flow and Delay* (McGraw-Hill, New York, 1964) (Reprinted by Dover Publications, New York, 1972).

[6] L. Kleinrock, *Queueing Systems, Vol. 1: Theory* (Wiley—Interscience, New York, 1975).

[7] L. Kleinrock, Performance of distributed multi-access computer communication systems, *Information Processing 1977, Proc. IFIP Congress*, Toronto, August 1977 (North-Holland, Amsterdam, 1977) 547–552.

[8] L. Kleinrock and F.A. Tobagi, Packet switching in radio channels: Part 1—Carrier sense multiple-access modes and their throughput-delay characteristics, *IEEE Trans. Commun.* **COM-23** (1975) 1440–1416.

[9] S. Lam, A new measure for characterizing data traffic, *IEEE Trans. Commun.* **COM-26** (1978) 137–140.

[10] J. Martin, *Systems Analysis for Data Transmission* (Prentice-Hall, Englewood Cliffs, NJ, 1972).

[11] B. Meister, J. Muller and H. Rudin, New optimization criteria for message-switching network, *IEEE Trans. Communication Technology* **COM-19** (1971) 256–260.

[12] L.G. Roberts, Dynamic allocation of satellite capacity through packet reservation, *AFIPS Conf. Proc., 1973 National Computer Conf., Vol. 42* (1973) 711–716.