LBL-13082
Preprint

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Engineering & Technical Services Division

To be published in EDP Performance Review

TOWARDS A HOLISTIC APPROACH TO NETWORK PERFORMANCE MANAGEMENT

D.F. Stevens

July 1981

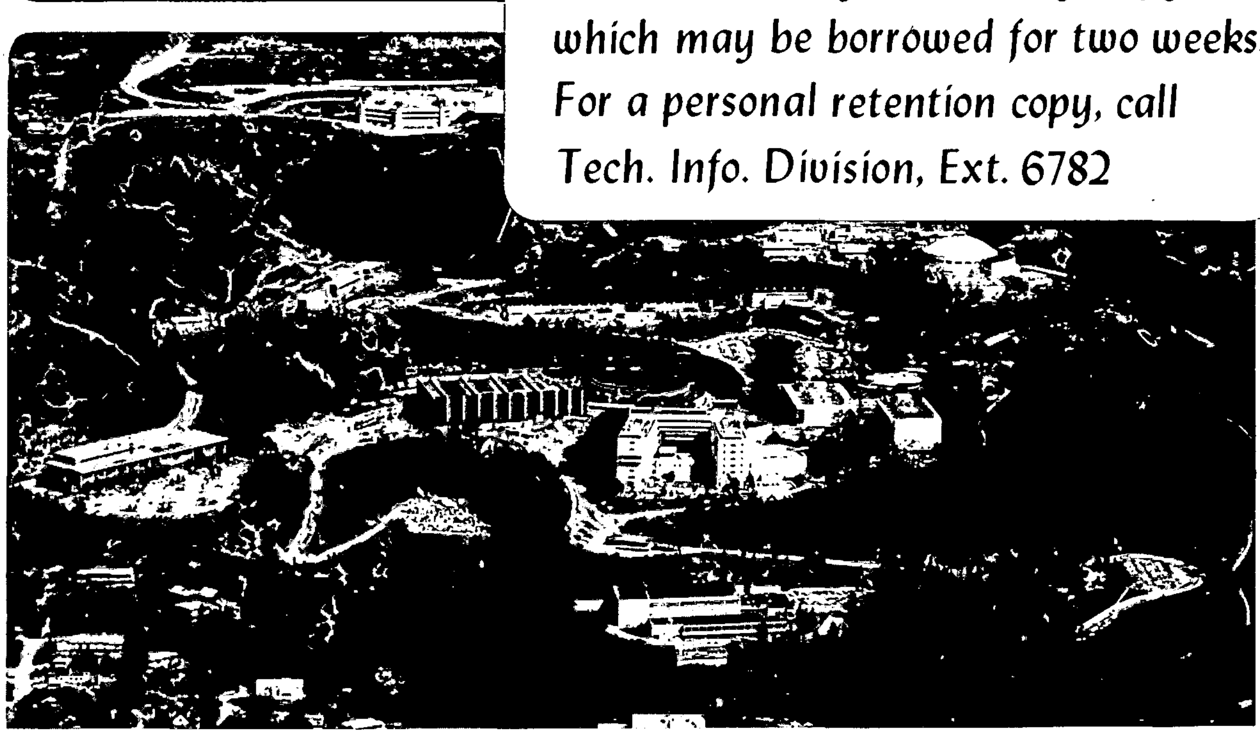

Prepared for the U.S. Department of Energy under Contract W-7405-ENG-48

To be published by
EDP Performance Review

Towards a Holistic Approach to Network Performance Management*

D. F. Stevens
Lawrence Berkeley Laboratory
University of California
One Cyclotron Road
Berkeley, California 94720

July 1981

Abstract: It is very easy to fall into the trap of managing a computer network as a collection of discrete components. The performance--and cost-performance--of a network, however, is not a simple function of the performance of its pieces. Familiar parameters--such as cost--are affected in unexpected ways by the presence of the network. Wholly new parameters are introduced. And management attitudes need re-evaluation in the light of network reality.

Keywords and Phrases: network(s), network management, network performance, network personality, network performance management, user-oriented management.

* This work was supported by the U.S. Department of Energy under contract No. W-7405-ENG-48.

Towards a Holistic Approach to Network Performance Management*

David F. Stevens
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

July 1981

holistic: adj. Emphasizing the importance of the whole and the interdependence of its parts.

The American Heritage Dictionary[1]

The whole is more than the sum of its parts.

G. Weinberg[2]

1. Introduction

A computer network is a poorly understood collection of (relatively) well-understood objects. One factor which contributes to our lack of understanding is our reluctance to look at the unfamiliar whole; our attention is caught by the familiar pieces--nodes and links--of which the whole is composed, and we are misled, as a result of our familiarity with elementary formal mathematics, into considering that the whole is precisely the sum of its parts. The real world is but an imperfect model for elementary formal mathematics, however, and Gerry Weinberg's observation (quoted above) is more likely than not to obtain in most real situations.

We here consider some of the ways in which it obtains for networks of com-

* This work was supported by the U. S. Department of Energy under contract No. W-7405-ENG-48.

puter systems. It is not my desire, in presenting this rather philosophical approach to network management, to persuade you to abandon the traditional technical concerns of computer management; rather, I hope to encourage you to complement those concerns with the holistic concerns introduced here. Furthermore, it is not my intent to tell you how to solve specific network problems but to present some ideas that encourage you to think in terms of the network as a whole. And finally, in the interest of brevity, I shall consider only three areas: the purpose and personality of networks, network performance parameters not found in stand-alone systems, and common pitfalls in network performance management.

2. The Purpose and Personality of Networks

Let us begin with my understanding of the fundamental purpose of a network and what it means to manage its performance:

> A network is a service entity existing only to connect users to hosts or to each other.
>
> Network performance management is the development and adjustment of the total system so as to achieve the kind (in numbers, power, availability, reliability, and usability) of connections the users want.

Network performance management should thus be a user-oriented activity. One of the clearly recognized consequences of the diffusion of the computer into our everyday lives is the rapid growth of the community of users. This growth includes not only more users, but also more kinds of users. In particular, it means more users by whom computer systems are--rightly--viewed in the same light as household appliances: they are labor-saving or convenience devices, the internal workings of which are of no interest, and the proper performance of which one ought to be able to take for granted. Some networks will be used almost exclusively by users with this point of view, some will retain a strongly technical user community, and some will have interesting blends of users. Proper performance management takes cognizance of the attitudes of the users and tailors the system accordingly.

One aspect of perceived performance which is particularly sensitive to the user blend, and yet which has been greatly neglected in discussions of performance management, is the personality of the system. The primary determinants of the personality of a network are the personalities of the constituent elements, the functional independence of the nodes, and the visible diversity of the links. The latter two are particularly important because they can either conceal or highlight differences in the individual elements. We can use them to define a two-dimensional "Personality Chart" for networks (see Figure 1): Region (a), with autonomous nodes and ad-hoc links, contains many home-grown networks which have been created by providing special, single-purpose communications paths to join pre-existing, independent, heterogeneous systems. Planned networks using differing computer architectures will tend to fall more in Region (b) (for despite the functional integration of the nodes the differing architectures give rise to multiple communication styles). The ARPANet--with its many completely independent nodes which adhere to a single communication protocol--is prototypical of Region (c), while Region (d) contains well-integrated, single-architecture, and often single-purpose, networks (such as some banking and airline reservation systems).

The apparent character of a network depends largely upon where it fits on the personality diagram. The further to the right it is the more a user will tend to see--and have to be aware of--the individual nodes; the further to the left, the more a user will tend to see a single system. The closer to the top it is the more a user will tend to see--and have to know--the precise protocol which allows Machine X to talk to Machine Y; the closer to the bottom it is, the more a user will tend to think in terms of network services, access to which is independent of the endpoints. (See Figure 2.) The difference between where your network fits on the Personality Chart and where your users want it to be should be one of the driving forces behind your management of the network.

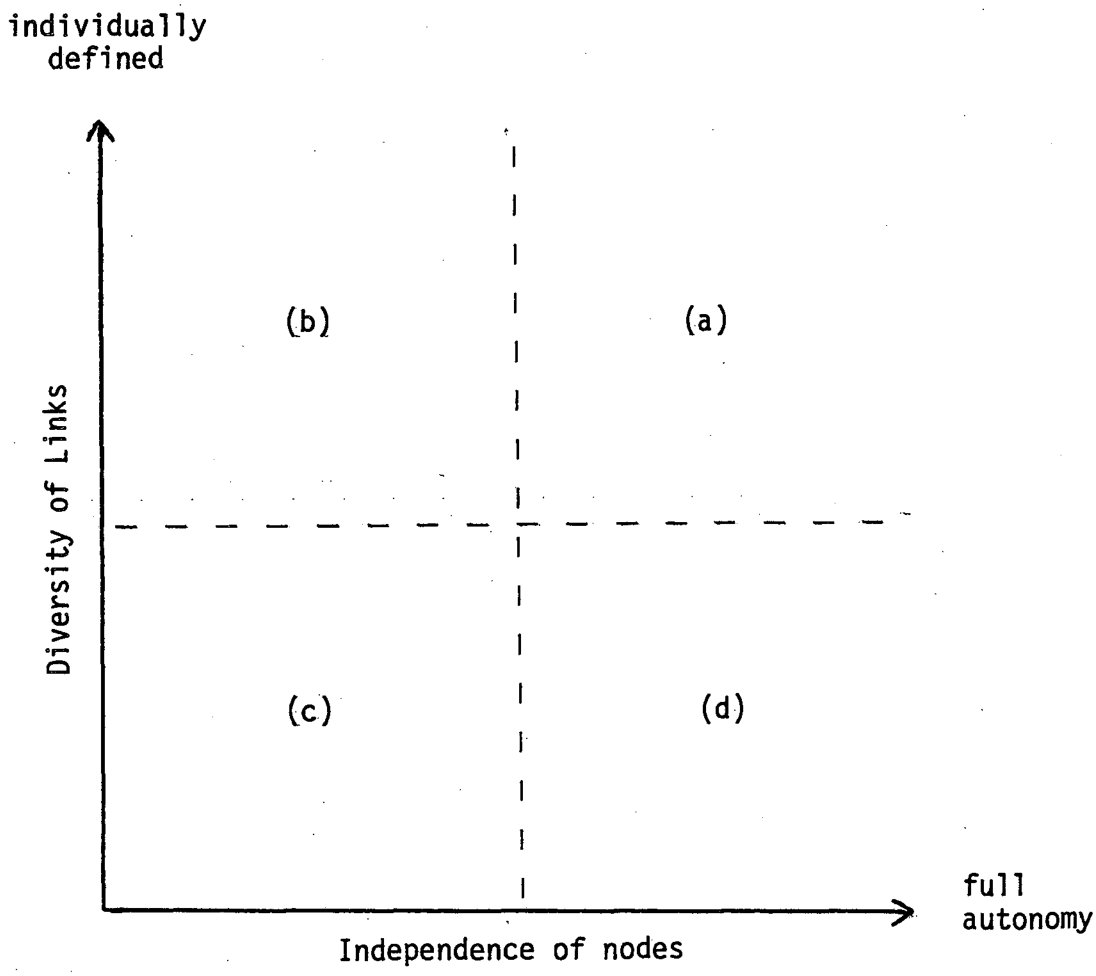

Figure 1: Network Personalities

a: autonomous nodes, ad-hoc links
b: functionally integrated nodes, ad-hoc links
c: autonomous nodes, single protocol
d: functionally integrated nodes, single protocol

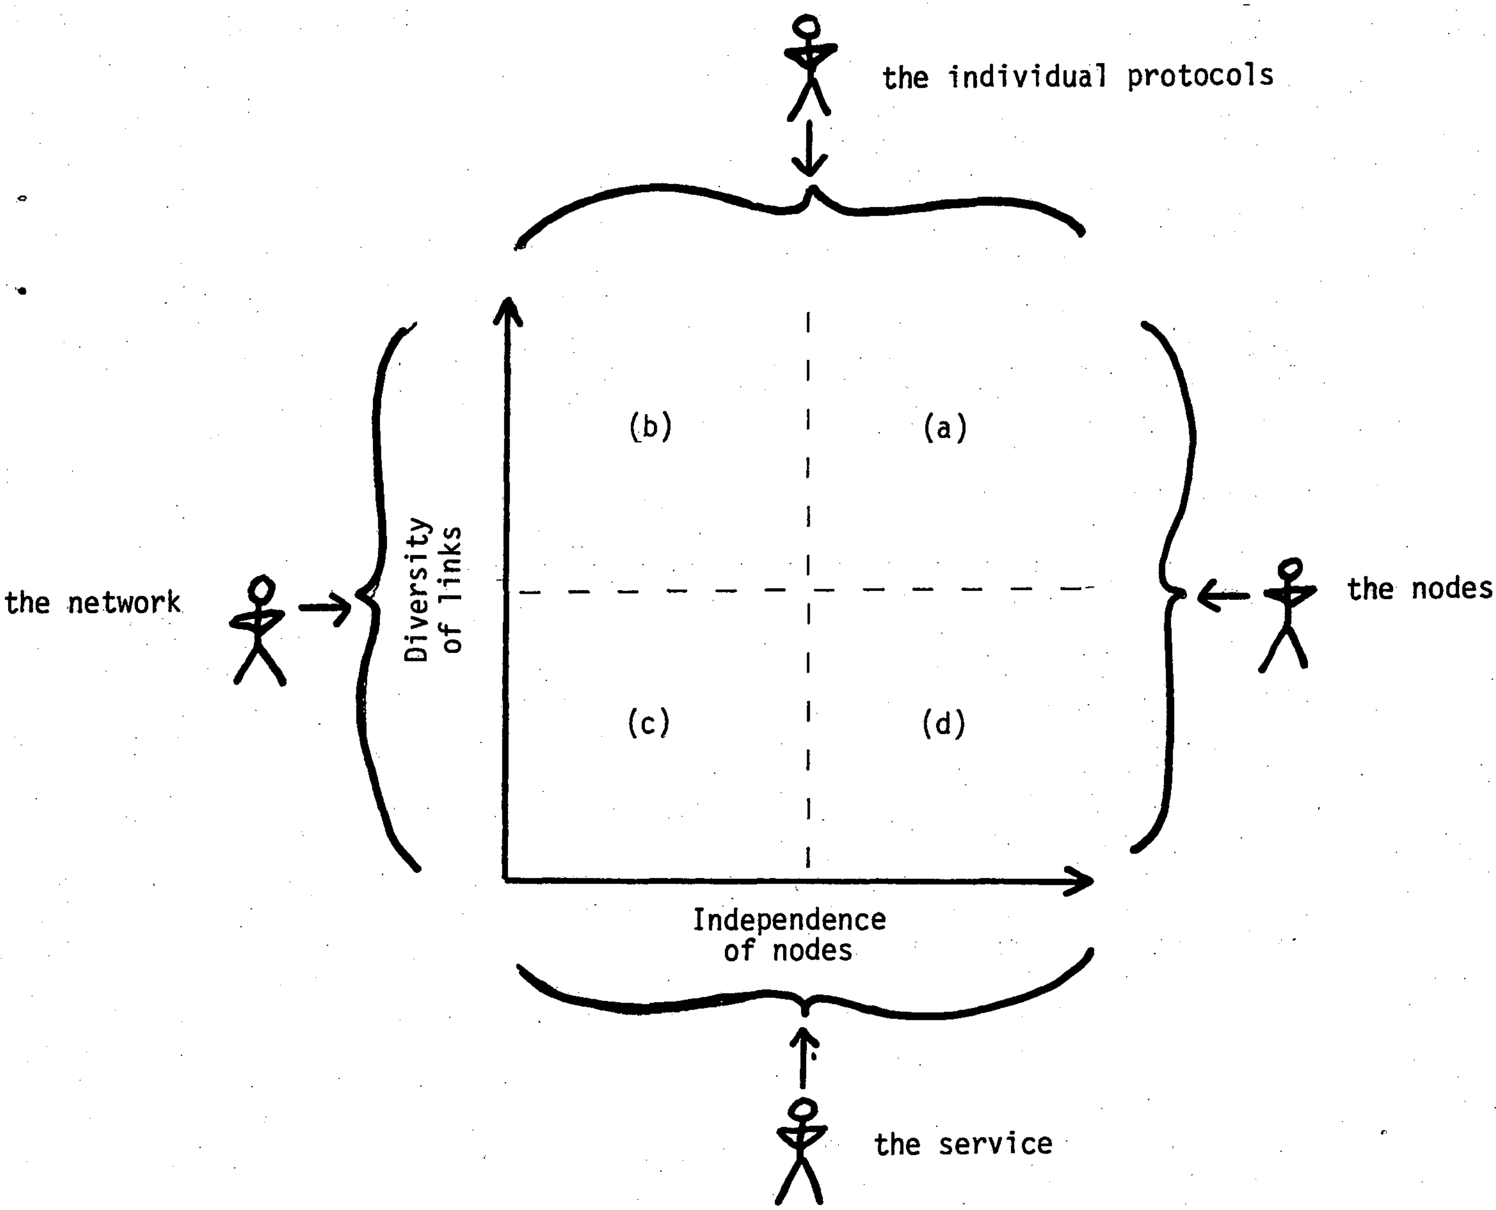

Figure 2: Network Personalities: The User's View

3. Performance Parameters Unique to Networks

Throughout its brief history the computer performance management game has been played on constantly shifting terrain: block-time gave way to batch use, sausage batch to multiprogramming, batch itself to interactive, discrete service to unified central service.... The newest shift sees centralized, stand-alone service giving way to distributed, network service. As with each earlier step, this new mode of operation brings a new environment, a new set of fundamental performance parameters. Three of these are discussed below. The first (cost of network support) is an extension of familiar concepts, the other two (topology and epidemiology) are quite new.

3.1 The cost of network support. The cost of network support exists in three levels of indirection. At Level 0 (not indirect at all) are the obvious hardware, software and communication costs; there is little point in discussing them here beyond noting that they exist. At Level 1 (partially indirect) there is the performance drain on each node. The indirection arises because active participation in the network is not a prerequisite for the accumulation of costs: mere connection is enough to cause a drain even if the node is a principal in none of the network traffic. In all cases there is some consumption of resources, even if only for residence of an unused network control program or device table; in some there can be significant amounts of network-induced overhead for such tasks as monitoring, listening, or forwarding.

The most significant costs are those at Level 2 (quite indirect); they manifest themselves as an increased complexity in the problem-handling process. This arises from many causes, the most obvious being that the system itself is more complex than in the pre-network days, and thus presents more complicated problems. But the very existence of a network tends to complicate the resolution of even simple problems: the distribution of computing is accompanied by a separation of symptoms and a dispersion of effects; it introduces conflicts of time, language, and custom; and it provides a much larger menu of possible causes for each symptom than formerly existed. These effects, in turn, lead to delays in the detection, diagnosis, and resolution of all problems, regardless of difficulty.

The detection of a problem is delayed because no local manager sees all the warning signs, because it may manifest itself in different ways at different sites, and because the language used to describe the same phenomenon may vary from site to site. After detection, diagnosis is slower, again because of language difficulties (not between natural languages--though they may exist and complicate things still further--but between the jargons used for describing system problems; these tend to be very architecture- and configuration-dependent and replete with historical references, metaphors, neologisms, and acronyms unique to each site) and because there are more possible sources of trouble. It can be quite difficult to distinguish faulty communication from sheer overloading, and either of those from design failure; it can be equally difficult to isolate terminal, host, node, and link problems one from another. Finally, after diagnosis, resolution is slower, again because of language difficulties (compounded by the fact that vendors are now involved and they often appear to have an unwritten policy to understand nothing not phrased in their own jargon), and also because of the coordination among several sites (with different priorities) that may be necessary.

3.2 Topology. The network manager needs to be aware of four aspects of generalized network topology: the physical topology, perceived access distance, dominance, and diversity. The first of these exhausts ones initial understanding of "topology"; it is tied very closely to the circuit diagram of the network, and includes such things as presence or absence of links or paths between two nodes, redundancy (the number and nature of alternative paths), path lengths (measured in links or seconds, but never in kilometers), and link capacity (which ought to be measured in user messages or user bits but usually isn't so restrictive). This view of topology is amply treated in more traditional approaches to network management.

Despite its name, perceived access distance has nothing to do with linear measure; it is a highly subjective impression of the remoteness of one site from another. (Suitably managed, two sites can seem to be quite far apart even if they are physically in the same room.) Among the factors helping to separate two sites are:

- response time, especially for simple requests
- differences in access protocols
- differences in terminology and phraseology
- differences in data formats
- error rates (reliable links are seen as short)

This aspect of topology is closely related to apparent diversity, of course. Diversity is a function of the number of different protocols a user must know and the number of kinds of nodes a user must understand to use the network effectively. (It increases as one moves North or East on the personality chart.) Note that it is not the absolute numbers which are emphasized, but how those numbers are reflected in the users' view of the network. Note further that there are two classes of users. The more numerous contains those who prefer to operate on a Principle of Least Knowledge: The less they have to know about technical computing details the better. Despite the frequent presence of high corporate officers in this first class, the other often has more influence on the style and personality of the network because it tends to be more vocal, to speak directly to the computer folk, and to speak in terms that computer folk appreciate: They want to know all about those technical details...and to be able to exploit their knowledge. Keeping both classes happy can be quite a trying experience.

The final aspect of topology in the network manager's burden is dominance. There are two forms of dominance, with very different characteristics. The more obvious sort is technical dominance, which results from a mismatch between the network's physical topology and technical capacity and its traffic pattern, to the extent that providing adequate service to the dominant links causes degradation of service elsewhere in the net. The other form of dominance is the political form, which produces a sort of tunnel vision in which only the "important" portions of the network retain visibility. The problem here, of course, is that small problems in the outlying portions of the net get no attention until they escalate or spread to such an extent that an "important" element is infected.

3.3 Epidemiology. Epidemiology is the study of the spread of disease. Its application to network management should be clear, but it seems to be largely ignored in the technical press. Epidemiology is less concerned with specific instances of network problems than with the following kinds of questions:

- What kinds of problem migrate or spread?
- Why? (What impels their movement?)
- How? (What medium carries them?)
- In what directions?
- How fast? How far?
- How can they be isolated? (How does one stop the spread, short of rupturing the network? If that's the only cure, when should it be applied?)
- How quickly?
- How can they be traced?

I have no answers for these questions; they--the answers--are largely net- and problem-specific anyway. But responsible network management demands that you, yourselves, pose them for your networks. The Great Power Blackouts in northeastern North America in 1965 and 1977 were epidemiological failures; do you wish to be responsible for what has been called "the Great Credit Blackout of 1990"[3]? (Smaller-scale contagions can be distressing, too: I'm sure most of you have--or know someone who has--experienced infectious disk contamination in the pre-winchester days.)

4. Pitfalls to be Avoided in Network Performance Management

Even an experienced DP manager--perhaps even _especially_ an experienced DP manager--can fall into habits of thinking which are less than optimal for network management. I will consider five of these traps for the unwary. They are all more concerned with point of view than with technique or specific knowledge.

4.1 Philosophy. It is essential that the network manager be aware of, and in sympathy with, the network philosophy. One must manage with the global view always in mind. Running one portion of a network at peak efficiency is injurious to the whole if it overloads other portions to the point of failure. (We've all experienced the results of expressway network management when the entrances are working more efficiently than the exits.) Do not manage primarily for capacity if availability is the design goal, or for response time if security is the more sensitive issue. Know the philosophy the network is intended to support and direct your efforts accordingly.

4.2 Locality. This trap is the hardest of all to avoid, especially for those experienced in the management of stand-alone systems. It manifests itself in many different ways, some of which are quite unexpected. Problems can be expected to arise with locality of time, of data, of definition, of language and custom, and of measurement.

Even the obvious problems with locality of time--for example, the existence in the contiguous United States of four time zones--are often forgotten in the convenience of modern telecommunication, with resultant misunderstandings. Far subtler, and far easier to forget, are the differences that can exist even within a single machine room, and which surface only as you try to reconstruct, from machine logs, the exact sequence of events which provoked a particular crisis. A different sort of locality of time--but just as awkward--is measured not in hours or milliseconds but in versions: the use of different versions of the same system at different nodes can yield quite interesting results.

There are two kinds of locality problems associated with data, physical and political. The physical problems are well-known, but no more tractable for their familiarity. Such questions as where data resides, or should reside; how many copies are kept, and where; which copy, or copies, should be updated, in what sequence, by whom, and when (according to which clock), are difficult enough on technical grounds, but can become nearly impossible when compounded with the political questions: Where does ownership and control reside? Which (whose) copy is to be believed when differences are detected?

Further complications can be introduced by non-uniform access procedures, which leads us into locality of language and custom. The effect on problem handling of local language and custom has already been noted; access is the other general area most affected by this sort of locality. It is a particularly difficult problem for networks formed by tying together a number of pre-existing stand-alone systems, but the Tower of Babel Principle ensures that strong forces exist to sow the seeds of diversity within even those networks designed and built, from the ground up, as networks.

Some of the problems caused by locality of definition are familiar to us in the story of the six blind men describing an elephant, or the boy who couldn't see the forest because no matter which way he looked, a tree was in the way. It leads us into error by encouraging inappropriate generalization and unjustifiable induction. One must always remember that a network has its own character, and that that character is often quite independent of--or at best obscurely related to--the characters of its constituent parts.

The fifth form of locality--locality of measurement--harks back to the opening quote from Gerry Weinberg. It is very tempting, for instance, to measure the individual performance of the several members of a network and to assert that the performance of the net is a vector, or product, or sum of those individual performance numbers. A fine example of this phenomenon appeared in the Spring, 1981, Computerworld/Extra! on Data Communications[4]. Although admiting that the approach was simplistic, the author defines "network availability" as

$$\prod_j \frac{\text{uptime}_j}{\text{uptime}_j + \text{downtime}_j} \left[= \prod_j \text{availability}_j \right]$$

where j ranges over all network elements. Leaving aside the considerable inadequacies of the implied definition of availability_j, this product form leads to meaningless results, as shown below:

Consider two small systems, one in series, the other in parallel:

I: 

II: 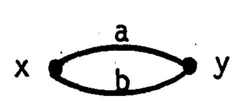

and assume for each system that "a" is up and "b" is down 0000-1200 each day, while "b" is up and "a" is down the other half of each day. Then the calculated "network availability" for both systems would be

$$Av = \frac{1}{2} \bullet \frac{1}{2} = \frac{1}{4}$$

but for any user trying to send a message from x to y, System I would always be down and System II would always be up.

To measure the performance of a network one must measure the performance of the network, not that of its components. (If one is lucky, it may turn out the one is derivable from the other, but the relations are usually quite diffuse, and definition via sums or products or vectors is rarely meaningful or justifiable.)

4.3 In(ter)dependence. The third potential trap for the network manager is failure to recognize the existence and strength of the interconnections among various parts of a network. The effects of interdependence are perhaps most easily recognized in instances of dominance, and one quickly learns how stresses in one part of a system affect the performance of other parts. The interdependence of reliability is more subtle, but some aspects--the propagation of electrical anomalies, for instance--are widely recognized, if not always completely understood. Some forms of data contamination can similarly propagate in mysterious ways to cause apparently unrelated failures in distant parts of the network.

The occasions that seem to have the most potential for disagreeable interdependence, however, are those involving unforeseen change. The problem can be most succinctly expressed by paraphrasing another observation of Gerry Weinberg: A network is a collection of parts no one of which can be changed[5]. In the urgency of the moment we are likely to forget that corrective action involves change, and so, all too often, in fixing one part of a network we debilitate, damage, or destroy another part.

4.4 Technicality. This is not a new problem, for informaticians have always had an over-fondness for the technical; nor is it unique to networks, for the seduction of technical measures has always lured the managers of performance away from their true responsibilities; but it seems to be a lesson that needs continuing reinforcement. The example given earlier of

$$\text{availability} = \prod_j \frac{\text{uptime}_j}{\text{uptime}_j + \text{downtime}_j}$$

serves equally well as an example of a technical measure which fails to represent in any way the users' view of the situation. Users are not interested in "uptime" and "downtime" (especially since much of the former occurs when they have no interest in the network): users care about "availability" with respect to their work. A much more user-oriented measure of availability would be

$$\text{availability}_i = \frac{\text{successful connections}}{\text{total attempts}}$$

where "successful connections" means that "the right message was received, by the right entity, in a timely manner, and understood without errors", and the subscript i indicates that you should know this number for each user. This is not a very technical measure, for many messages are technically correct but fail one or more of the criteria for "successful connections", but it will do more to tell you what your users think of your network than any technical measure will. If you wish to manage your network for the benefit of your users you must supplement your technical measures with measures that are applicable, appropriate, and user-oriented.

4.5 Purpose. The final network management trap I wish to bring to your attention is that of believing that actual use of the network matches its announced purpose. This, too, may be succinctly expressed by paraphrasing a systems theorist (in this case, John Gall): Networks develop goals of their own the instant they come into being[6]. Perhaps the most well-known instance of this is the ARPANet, which was created to do computer science research but which continues to exist as one of the world's most expensive memo distribution services (more than 60% of the total traffic is reputed to be electronic mail).

It has always been the case in computing that use outstrips purpose. Simultaneous with the development of networks is a tremendous expansion of the user community, in both numbers and breadth. The more people with fewer well-developed notions about the limitations of computing systems that enter the user community, the more the spectrum of uses to which the system is bent will widen, and the more unofficial and unanticipated network goals will develop.

5. Conclusion

(With apologies to John Donne, John Gall, and Gerry Weinberg ([7], [8], annd [2], respectively).) No node is an island sufficient unto itself alone. In managing networks it is important to avoid the simplifications achievable by dealing only with individual components: "success" in any single node may be failure in the network to which it belongs. A network is more than the sum of its parts.

References

[1] AHD, New College Edition, Houghton Mifflin, 1976, p. 628.

[2] G. Weinberg, An Introduction to General Systems Thinking, Wiley, 1975, p. 43.

[3] Oral comment in Session 52B, IFIP-77.

[4] F. Dzubeck, Network Management, Computerworld/Extra!, April, 1981, p. 91.

[5] G. Weinberg, op. cit., p. 162.

[6] J. Gall, Systemantics, Quandrangle, 1977, p.59.

[7] J. Donne, Devotions (#XIV, Nunc lento sonitu dicunt, Morieris).

[8] J. Gall, op. cit., p. 92.

TECHNICAL INFORMATION DEPARTMENT
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720