A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at:
www.cisco.com/ipj

F R O M   T H E   E D I T O R

Numerous technologies have been developed to protect or isolate corporate networks from the Internet at large. These solutions incorporate security, either end-to-end (IP security, or IPSec), or at the Internet/intranet border (firewalls). A third class of systems allow a range of IP addresses to be used internally in a corporate network, while preserving IP address consumption through the use of a *single* public address. This latter class of device is called a *Network Address Translator* (NAT), and while many Internet engineers consider NATs to be "evil," they are nonetheless very popular. Combining IPSec, NATs, and firewalls can be quite challenging, however. In our first article Lisa Phifer explains the problem and offers some solutions.

Successful network design is the result of many factors. In addition to the basic building blocks of routers, switches and circuits, network planners must carefully consider how these elements are interconnected to form an overall system with as few single points of failure as possible. In our second article, Valdis Krebs looks at how lessons learned from social network analysis can be applied to the design of computer networks.

The current Internet grew out of several government-funded research efforts that began in the late 1960s. Today, basic technology development as well as research into new uses of computer networks continues in many research "testbeds" all over the world. Bob Aiken describes the past, present and future state of network research and research networks.

The online subscription system for this journal will be up and running in January at **www.cisco.com/ipj.** In addition to offering a subscription form, the system will allow you to select delivery options, update your mailing and e-mail address, and much more. Please visit our Web site and give it a try. If you encounter any difficulties, please send your comments to **ipj@cisco.com.**

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

The Social Life of Routers
Applying Knowledge of Human Networks to the Design of Computer Networks

by Valdis Krebs

We often forget that computer networks are put in place to support human networks—person-to-person exchanges of information, knowledge, ideas, opinions, insights, and advice. This article looks at a technology that was developed to map and measure human networks—social network analysis—and applies some of its principles and algorithms to designing computer networks. And as we see more peer-to-peer (P2P) models of computer-based networks, the P2P metrics in human network analysis become even more applicable.

Social network analysts look at complex human systems as an interconnected system of nodes (people and groups) and ties (relationships and flows)—much like an internetwork of routers and links. Human networks are often unplanned, emergent systems. Their growth is sporadic and self-organizing[1]. Network ties end up being unevenly distributed, with some areas of the network having a high density of links and other areas of the network sparsely connected. These are called "small world networks"[2]. Computer networks often end up with similar patterns of connections—dense interconnectivity within subnetworks, and sparser connections uniting subnetworks into a larger internetwork.

Social network researchers and consultants focus on geodesics—shortest paths in the network. Many of today's social network algorithms are based on a branch of mathematics called graph theory. Social network scientists have concentrated their work, and therefore their algorithms, in the following areas:

• Individual node centrality within a larger network—network dependency and load upon individual routers

• Overall path distribution—good connectivity without excessive routing tables

• Improving communication flow within and between groups—designing better topologies

• Network patterns surrounding ego networks—strategies for analyzing and manipulating individual router connections

• Analyzing information flow behavior of client organization—how computer networks can support human networks

One of the methods used to understand networks and their participants is to evaluate the location of actors in the network. Measuring the network location is finding the *centrality* of a node[3]. All network measures discussed here are based on geodesics—the shortest path between any two nodes. We will look at a social network, called the *kite network*, that effectively shows the distinction between the three most popular centrality measures—the ABCs—*A*ctivity, *B*etweenness, and *C*loseness.

This model[4] was first developed by David Krackhardt, a leading researcher in social networks.

Activity

Figure 1 shows a simple social network. A link between a pair of nodes depicts a bidirectional information flow or knowledge exchange between two individuals. Social network researchers measure network activity for a node by using the concept of degrees—the number of direct connections a node has.

In this human network, Diane has the most direct connections in the network, making hers the most active node in the network with the highest degree count. Common wisdom in personal networks is "the more connections, the better." This is not always so. What really matters is where those connections *lead to*—and how they connect the otherwise unconnected![5] Here Diane has connections only to others in her immediate cluster—her clique. She connects only those who are already connected to each other—does she have too many redundant links?

Figure 1: Human Network
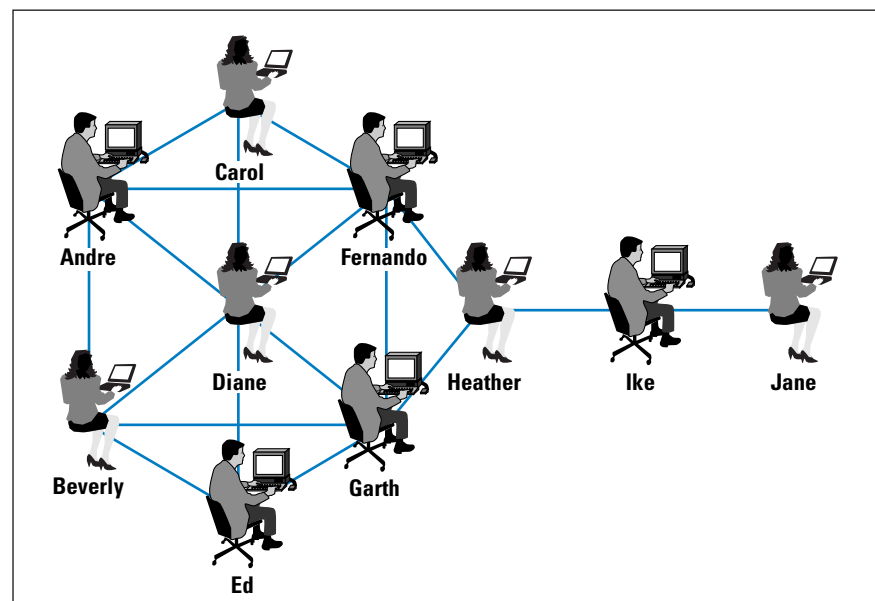
Betweenness

While Diane has many direct ties, Heather has few direct connections—fewer than the average in the network. Yet, in may ways, she has one of the best locations in the network—she is a boundary spanner and plays the role of broker. She is *between* two important constituencies, in a role similar to that of a border router. The good news is that she plays a powerful role in the network, the bad news is that she is a single point of failure. Without her, Ike and Jane would be cut off from information and knowledge in Diane's cluster.

Closeness

Fernando and Garth have fewer connections than Diane, yet the pattern of their ties allow them to *access* all the nodes in the network more quickly than anyone else. They have the shortest paths to all others—they are *close* to everyone else. Maximizing closeness between *all* routers improves updating and minimizes hop counts. Maximizing the closeness of only one or a few routers leads to counterproductive results, as we will examine below.

Their position demonstrates that when it comes to network connections, quality beats out quantity. Location, location, location—the golden rule of real estate also works in networks. In real estate it is geography—your physical neighborhood. In networks, it is your virtual location determined by your network connections—your network neighborhood.

Network Centralization

Individual network centralities provide insight into the individual's location in the network. The relationship between the centralities of all nodes can reveal much about the overall network structure. A very centralized network is dominated by one or a few very central nodes. If these nodes are removed or damaged, the network quickly fragments into unconnected subnetworks. Highly central nodes can become critical points of failure. A network with a low centralization score is not dominated by one or a few nodes—such a network has no single points of failure. It is resilient in the face of many local failures. Many nodes or links can fail while allowing the remaining nodes to still reach each other over new paths.

Average Path Length in Network

The shorter the path, the fewer hops/steps it takes to go from one node to another. In human networks, short paths imply quicker communication with less distortion. In computer networks, the signal degradation and delay is usually not an issue. Nonetheless, a network with many short paths connecting all nodes will be more efficient in passing data and reconfiguring after a topology change.

Average Path Length is strongly correlated with Closeness throughout the network. As the closeness of all nodes to each other improves (average closeness), the average path length in the network also improves.

Internetwork Topology

In the recent network design book, *Advanced IP Network Design* [6], the authors define a well-designed topology as the basis of a well-behaved and stable network. They further propose that "three competing goals must be balanced for good network design":

- Reducing hop count
- Reducing available paths
- Increasing the number of failures the network can withstand

Our social network algorithms can assist in measuring and meeting all three goals.

- Reducing the hop count infers minimizing the average path length throughout the network—maximize the closeness of all nodes to each other.
- Reducing the available paths leads to minimizing the number of geodesics throughout the network.
- Increasing the number of failures a network can withstand focuses on minimizing the centralization of the whole network.

On the following pages we examine various network topologies and evaluate them using social network measures while remembering these three competing goals of network design.

The models we examine do *not* cover hierarchical structures—with Core, Distribution, and Access layers—found in networks of hundreds or thousands of routers. We examine flat, nonhierarchical topologies such as those found in smaller internetworks, area subnetworks, or within core backbones. The topologies we model are the most commonly used—Star, Ring, Full Mesh, and Partial Mesh. We compute the social network measures on each of the topologies and discuss how the various measures help us meet the competing goals discussed above.

Star Topology

The Star topology, shown in Figure 2, has many advantages—but one glaring fault. The advantages include ease of management and configuration for the network administrators. For the Star, the three competing goals delineate as follows:
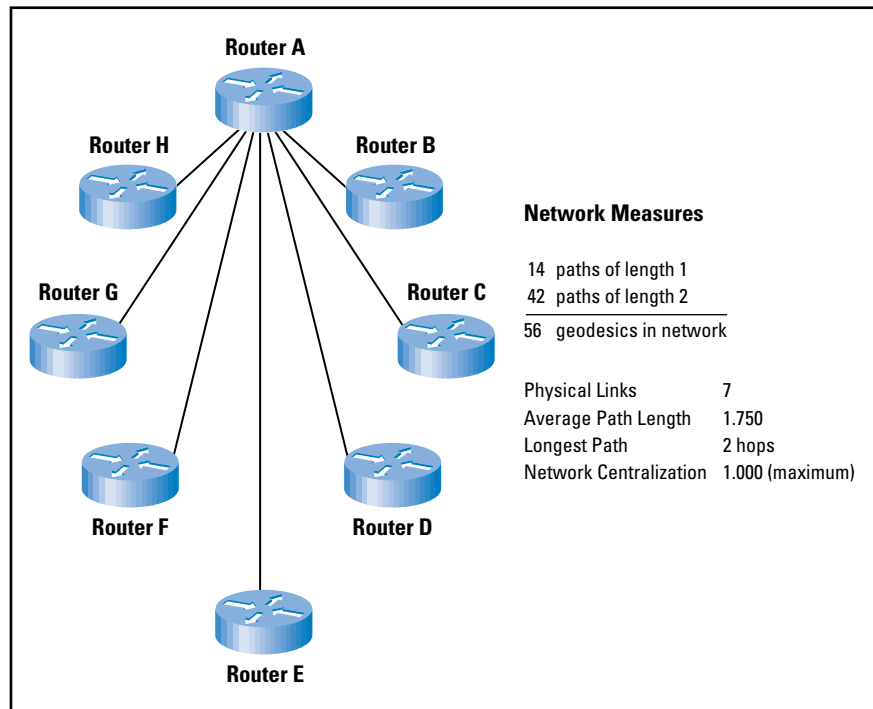
- *Reducing hop count:* The short average path length (1.75) throughout the network meets this goal well. Any router can *reach* any other router in two steps or less.
- *Reducing available paths:* The fact that there are a minimum number of possible available paths (56) to reach all other nodes—will not overload the routing tables, nor cause delays during routing table updates. It takes only seven bidirectional links to create the available paths.

• *Reducing network failures:* The network fails miserably if Router A goes down. Also, any link failure isolates the attached router—there are no multiple paths to reach each router.

Router A is not only a single point of failure—it is also a potential bottleneck—it will likely become overburdened with packet flows and routing updates as more routers are added in the star structure.

Router A receives the top score (1.000) in Activity, Betweenness, and Closeness. As a result, the network is very centralized around Router A from the perspective of all measures.

Figure 2: Routers in Star Topology

Network Measures

14 paths of length 1
42 paths of length 2
_____
56 geodesics in network

Physical Links   7
Average Path Length   1.750
Longest Path   2 hops
Network Centralization   1.000 (maximum)
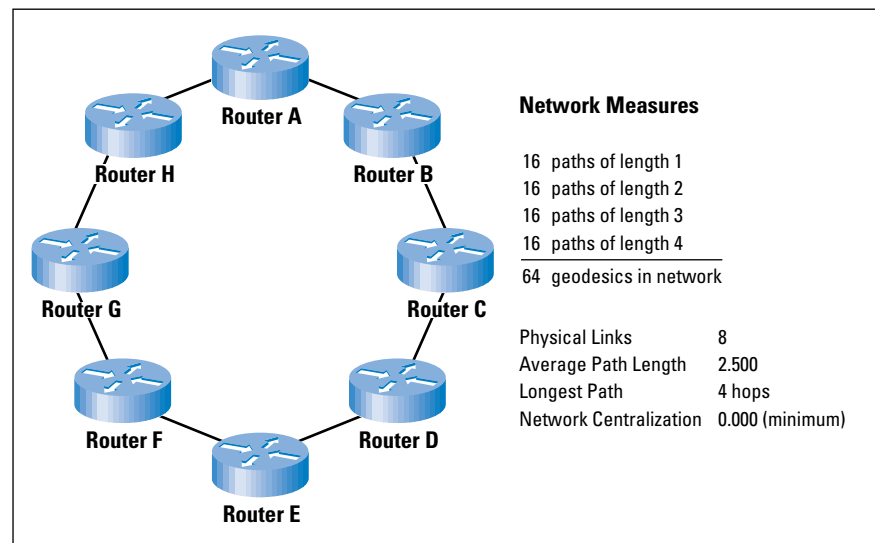
Ring Topology

The Ring topology, shown in Figure 3, is an improvement over the Star. It has some of the same advantages, but does not eliminate all of the drawbacks of the Star. The advantages include ease of management and configuration for the network administrators—adding another router is very simple. Unlike the Star topology, the Ring provides some redundancy and, therefore, eliminates the single point of failure—all nodes have an alternate path through which they can be reached. Yet it is still vulnerable to both link and router failures. For the Ring, the three competing goals delineate as follows:

• *Reducing hop count:* The average path length of 2.5 is quite long for a small network of eight nodes. Some routers (that is, A and E) require four steps to reach each other! Many ring physical layers hide this complexity from the IP layers in order to make those hops invisible to routing protocols.

- *Reducing available paths:* This configuration has more geodesics (64) than Star, yet not significantly more to overload the routing tables, nor cause delays during table updates.
- *Reducing network failures:* Even though network centralization is at the minimum (no node is more central than any other), this network reaches failure quickly because of its weak redundancy. The Ring topology can withstand one link failure or one router failure and still keep a contiguous network. Two simultaneous failures can cause unreachable segments because of the lack of redundancy.

Most modern ring technologies such as *Synchronous Optical Network* (SONET) or the Cisco *Dynamic Packet Transport Protocol* (DPT) add a measure of redundancy by running a dual ring that heals itself if a link gets cut. The network "wraps" to avoid the downed line and operates at lower speed. A two-hop path can become a six-hop path if a single link fails. This can cause network congestion if the original dual ring was being used for data in all directions.

Figure 3: Routers in Ring Topology

Network Measures

16 paths of length 1
16 paths of length 2
16 paths of length 3
16 paths of length 4

64 geodesics in network

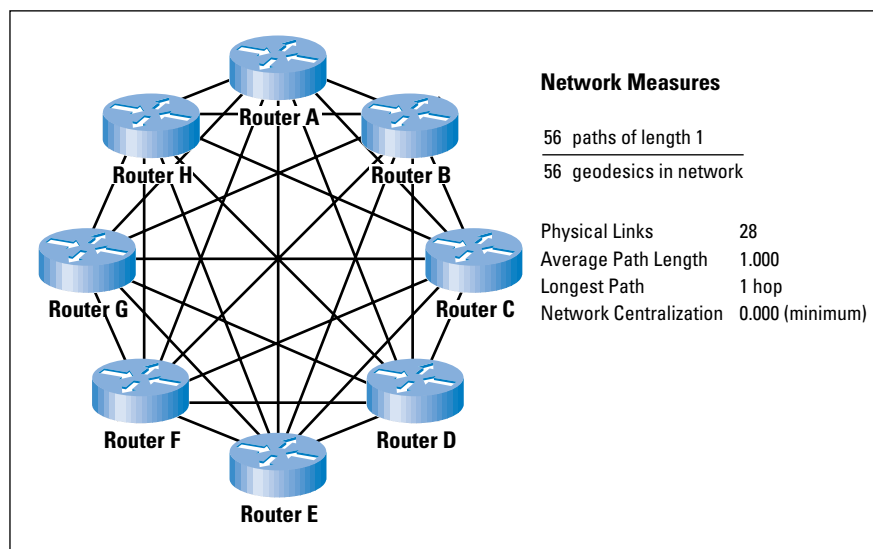| Physical Links	8
Average Path Length	2.500
Longest Path	4 hops
Network Centralization	0.000 (minimum)

Full Mesh Topology

The Full Mesh topology has several big advantages and several faults. The advantages include short path length (one hop) to all other routers and maximum resilience to failure if links or routers start failing. The disadvantages revolve around the complexity created by this topology. For the Full Mesh, the three competing goals delineate as follows:
- *Reducing hop count:* The shortest path length possible is attained for all routes—all nodes can reach each other in one hop.
- *Reducing available paths:* There are a minimum number of possible available paths (56) to reach all other nodes. The routing entries will not overload the routing tables, nor cause delays during routing table updates.

• *Reducing network failures:* The network is not dependent upon any single node (network centralization = 0.000). This configuration represents the most robust topology available—chances are very slim that the number of failures necessary to fragment the network will actually occur within the same time period.

The disadvantages of the Full Mesh topology all focus on one glaring fault—there are too many physical links. If the routers are far apart, the link costs can quickly become prohibitively expensive because adding routers creates a geometrical explosion in links required—soon the routers do not have enough ports to support this topology. Administering the system and keeping an up-to-date topology map becomes more and more complex as routers are added. The network in Figure 4 has 28 two-way links. Double the routers, in a full mesh topology, and the link count increases by a factor greater than 4.

Figure 4: Routers in Full Mesh Topology

Network Measures

56	paths of length 1
56	geodesics in network

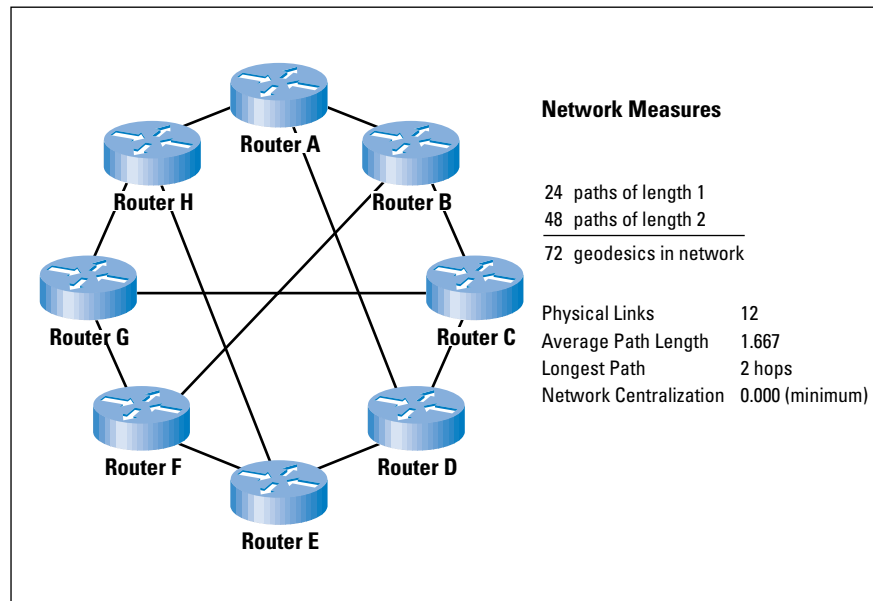|	
Physical Links	28
Average Path Length	1.000
Longest Path	1 hop
Network Centralization	0.000 (minimum)

Partial Mesh Topology

The Partial Mesh topology is quite different. It is the most difficult to build—there is no simple rule to follow (rule for Star: connect everyone to Router A; rule for Full Mesh: connect everyone to everyone). If built incorrectly, the partial mesh layout can have many of the disadvantages of the former topologies without many of the benefits. If built correctly, the opposite is true—more advantages, fewer disadvantages.

Building a successful partial mesh topology is where the interactive use of our social network measures really comes into play. The design below evolved after several iterations. With every iteration the average path length dropped until it appeared to reach a plateau where no further changes lowered the hop count without noticeably increasing the number of physical links. For the Partial Mesh, the three competing goals delineate as follows:

- *Reducing hop count:* The short average path length (1.667) throughout the network meets this goal well. Any router can *reach* any other router in two steps or less. Path length is less than that for the Star and Ring topologies.
- *Reducing available paths:* The number of available paths in the network (72) is the highest among all topologies, though not significantly more than the Ring topology. As the number of nodes in a network increases, this could become a problem—the average path length vs. path count trade-off needs to be closely monitored.
- *Reducing network failures:* Network centralization (0.000) is the same as for the Full Mesh topology—no router, nor link, is more important than any other. As nodes or links are removed from this network, it does not fragment quickly. Chances are slim that the number of failures necessary to fragment the network will actually occur within the same time period. Although we optimized our network centralization for this small "toy" network, we cannot expect this for most real networks. Yet, the goal remains to keep this metric as small as possible.

This topology in Figure 5 was built starting with a Ring topology—a simple architecture. A link was added and the network was remeasured. Was this structure better than the previous? If so, the current structure was kept and another link was added and the network was remeasured. This iterative process was continued until no further improvements happened after several changes. This process does not guarantee an *optimum* solution, yet it quickly converges on a *good* solution—even large networks improve quickly with just a few added links.

Figure 5: Routers in Partial Mesh Topology

Network Measures

24 paths of length 1
48 paths of length 2
────────────────
72 geodesics in network

Physical Links   12
Average Path Length   1.667
Longest Path   2 hops
Network Centralization   0.000 (minimum)

A quirky aspect of networks is that sometimes you can subtract by adding—add a link to a network and reduce the average path length. The opposite also works, sometimes. You can add by subtracting—remove a tie and watch the average hop count grow. Yet, you never know for certain what effect adding or removing a link will have—it is neither a linear nor a local phenomenon. The size and direction of these changes depend upon the existing topology of the network and the location of the added or removed tie. It is key to have a model that allows quick what-if calculations.

Let's experiment with removing random ties—a situation similar to links between routers failing. If we remove the link between Router A and Router H in Figure 5, the number of geodesics in the network increases from 72 to 76, and the average path length increases to 1.815. Yet, removing a different link, G to F, reduces the the number of geodesics in the network from 72 to 66, while the average path length increases only to 1.727. If we are concerned about too many paths in the network, we can remove another link, B to C. This further decreases the number of shortest paths to 60, while reducing physical links to 10. This is very near the 56 paths in the very efficient star topology. Whereas the star is very vulnerable because of its single point of failure, this partial mesh, with the two links removed, is still robust. While the number of geodesics drops, the average path length creeps up slightly to 1.80 with the removal of the second link. Figure 5 has no paths greater than two hops. With the two links (G to F, B to C) removed, we now have 8 geodesics of three hops, while at the same time 12 fewer geodesics to load into routing tables, and two fewer physical links. It is a constant trade-off.

NSFnet Backbone
The NSFnet Backbone network, shown in Figure 6, connected the supercomputing centers in the USA in 1989. It is a partial mesh design that functions as a real-life example to test our social network algorithms.

Figure 6: NSFnet in 1989
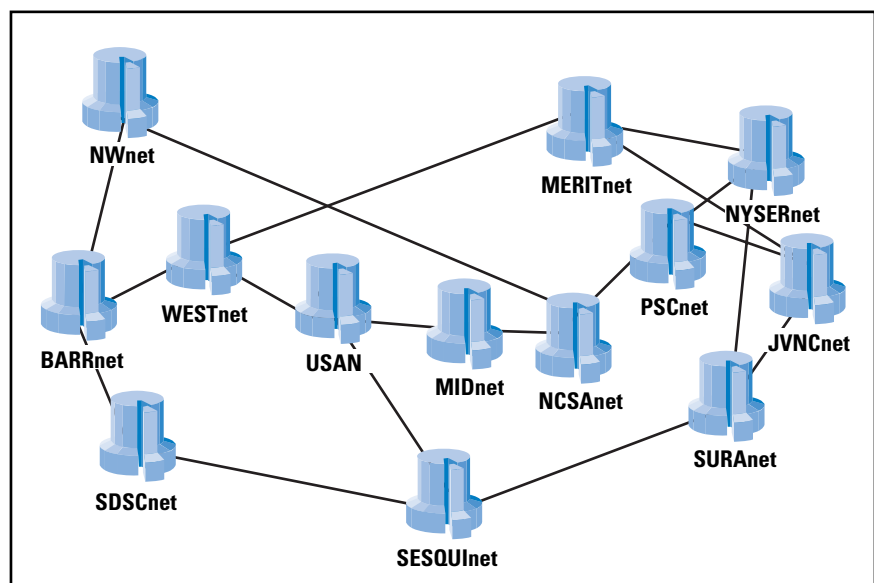
We remember our three competing goals for good internetwork design.

- Reducing hop count: average path length in steps/hops
- Reducing available paths: total geodesics in the network
- Increasing the number of failures the network can withstand: network centralization

What happens to these goals as we experience failures in the links or the nodes of the network? Table 1 shows the base metrics for Figure 6 and then shows what happens to the metrics, and our three goals, when five different failures occur.

Table 1: Possible Link and Node Failures

Scenario	Number of Geodesics in the Network	Network Centralization	Longest Path (hops)	Average Path Length (hops)
Original Design (Figure 6)	200	0.062	4	2.370
1) Node failure: NCSA	180	0.208	5	2.689
2) Node failure: MID	180	0.083	4	2.489
3) Node failure: JVNC	148	0.046	4	2.324
4) Link failure: NCSA–PSC	230	0.167	6	2.974
5) Link failure: USAN–MID	212	0.123	5	2.660
6) Link failure: MERIT–JVNC	192	0.069	4	2.458

The most damaging was link failure 4—the link failure between NCSA and PSC. This link is between two of the most central nodes in the network. If the flows between nodes are distributed somewhat evenly, then this link is one of the most traveled in the network.

The least damaging is node failure 3—the node failure at JVNC. In fact, this failure improved most metrics! By removing this node from the network, the number of network paths drops significantly, network centralization decreases, path length decreases slightly, and the longest path is still four hops.

The original NSFnet topology design is very efficient. I tried two different strategies to improve the network. The first strategy involved moving existing links to connect different pairs of routers. No obviously better topology was found by rearranging links among the routers. I was not able to find a better design that reduced both the number of geodesics and the average path length without significantly increasing the number of physical links in the network.

The second strategy is counter-intuitive, yet often networks respond well to this approach. It is the "subtracting by adding" approach described above. By adding new links in the right place in the network, we not only reduce the distance between nodes, we also decrease the number of geodesics in the network.

Because the NSFnet nodes had a maximum limit of three direct neighbors, I started connecting the nodes of Degree = 2. Options 1 through 3 show the various combinations and their effect on the total network. The improvements are minimal, yet each option offers specific strengths.

Option 2 offers more improvements than the others.
- The longest geodesic was reduced to three hops.
- The average path length was reduced throughout the network.
- The number of paths for the routers to remember was reduced slightly.
- Network centralization did not increase enough to noticeably affect the number of failures the network could withstand.

Table 2: Possible Network Improvements

Scenario	Number of Geodesics in the Network	Network Centralization	Longest Path (hops)	Average Path Length (hops)
Original Design (Figure 6)	200	0.062	4	2.370
Option 1 (add link: SDSC–MID)	202	0.071	4	2.287
Option 2 (add link: NW–DSC)	198	0.074	3	2.273
Option 3 (add link: NW–MID)	202	0.050	4	2.356

The improvement in Option 2 (add link: NW–SDSC) was actually implemented in the 1991 version of NSFnet—an excellent example of the "subtracting by adding" network dynamic. Networks are complex systems. How the network responds to change is based on the distribution and pattern of connections throughout the network.

Conclusion

In the real world we may not have the flexibility to experiment with our network model as we have with these examples. There will be more constraints. The information flows in your organization may require that specific pairs of routers have direct links—even if those connections would not be recommended by the algorithms we have been examining. Yet, when we have our "must-have" connections in place, we can experiment with the placement of the remaining connections using these social network metrics to indicate when we are getting close to a robust, yet efficient topology.

Given "initial conditions," social network methods can model our computer networks and suggest link changes[7] to form an effective topology that has a short average hop count, not too many paths, and just enough redundancy.

References

[1] Krebs V., "Visualizing Human Networks," *Release 1.0*, Esther Dyson's Monthly Report, February 1996.

[2] Watts D., Strogatz S., "Collective Dynamics of Small World Networks," *Nature*, 4 June 1998.

[3] Freeman L., "Centrality in Social Networks: A Conceptual Clarification," *Social Networks*, No. 1, 1979.

[4] Krackhardt D., "Assessing the Political Landscape: Structure, Cognition, and Power in Organizations," *Administrative Science Quarterly*, No. 35, 1990, page 351.

[5] Burt, Ronald S., *Structural Holes—The Social Structure of Competition*, ISBN 0674843711, Harvard University Press, 1992.

[6] Retana, A., Slice, D., White, R., *Advanced IP Network Design*, ISBN 1578700973, Cisco Press, 1999.

[7] Hagen G., Discussions with fellow network researcher, Guy Hagen, regarding combinatorial algorithms and models for recommending changes to improve the overall topology of a network. Link to web site.

VALDIS E. KREBS leads his own management consulting firm—orgnet.com He holds an undergraduate degree in Mathematics & Computer Science and a graduate degree in Human Resources. Since 1988 he has applied organizational network analysis to improve knowledge work within and between Fortune 500 firms such as IBM, Lucent, TRW, and supported consulting firms such as Ernst & Young, PricewaterhouseCoopers, and Booz-Allen-Hamilton. In addition to knowledge networks, he has applied these methodologies to mapping, measuring, and molding strategic alliances, communities of interest, emergent structures on the WWW, and internetworks. His work has been referenced in many publications, including the *Wall Street Journal*, *Entrepreneur*, *Training*, *PC Magazine*, *ZDNet*, *Corporate Leadership Council's Best Practices Reports*, *Knowledge Management*, *Across the Board*, *Business Week*, *HR Executive*, *Personnel Journal*, *FORTUNE*, and Esther Dyson's influential information industry newsletter, *Release 1.0*. He writes a regular column, "Working in the Connected World," for the *IHRIM Journal*. His Web site is at: www.orgnet.com and his e-mail is: valdis@orgnet.com