Survivability of Multiagent-Based Supply Networks: A Topological Perspective

Hari Prasad Thadakamalla, Usha Nandini Raghavan, Soundar Kumara, and Réka Albert, *Pennsylvania State University*

You can improve a multiagent-based supply network's survivability by concentrating on the topology and its interplay with functionalities.

Supply chains involve complex webs of interactions among suppliers, manufacturers, distributors, third-party logistics providers, retailers, and customers. Although fairly simple business processes govern these individual entities, real-time capabilities and global Internet connectivity make today's supply chains complex. Fluctuating demand patterns, increasing customer expectations, and competitive markets also add to their complexity.

Supply networks are usually modeled as multiagent systems (MASs).[1] Because supply chain management must effectively coordinate among many different entities, a multiagent modeling framework based on explicit communication between these entities is a natural choice.[1] Furthermore, we can represent these multiagent systems as a complex network with entities as nodes and the interactions between them as edges. Here we explore the survivability (and hence dependability) of these MASs from the view of these complex supply networks.

Today's supply networks aren't dependable—or survivable—in chaotic environments. For example, Figure 1 shows how mediocre a typical supply network's reaction to a node or edge failure is compared to a network with built-in redundancy.

Survivability is a critical factor in supply network design. Specifically, supply networks in dynamic environments, such as military supply chains during wartime, must be designed more for survivability than for cost effectiveness. The more survivable a network is, the more dependable it will be.

We present a methodology for building survivable large-scale supply network topologies that can extend to other large-scale MASs. Building survivable topologies alone doesn't, however, make an MAS dependable. To create survivable—and hence dependable—multiagent systems, we must also consider the interplay between network topology and node functionalities.

A topological perspective

To date, the survivability literature has emphasized network functionalities rather than topology. To be survivable, a supply network must adapt to a dynamic environment, withstand failures, and be flexible and highly responsive. These characteristics depend on not only node functionality but also the topology in which nodes operate.

The components of survivability

From a topological perspective, the following properties encompass survivability, and we denote them as survivability *components*.

The first is *robustness*. A robust network can sustain the loss of some of its structure or functionalities and maintain connectedness under node failures, whether the failure is random or is a targeted attack. We measure robustness as the size of the network's largest

1541-1672/04/$20.00 © 2004 IEEE
Published by the IEEE Computer Society

connected component, in which a path exists between any pair of nodes in that component.

The second is *responsiveness*. A responsive network provides timely services and effective navigation. Low characteristic path length (the average of the shortest path lengths from each node to every other node) leads to better responsiveness, which determines how quickly commodities or information proliferate throughout the network.

The third is *flexibility*. This property depends on the presence of alternate paths. Good clustering properties ensure alternate paths to facilitate dynamic rerouting. The clustering coefficient, defined as the ratio between the number of edges among a node's first neighbors and the total possible number of edges between them, characterizes the local order in a node's neighborhood.

The fourth is *adaptivity*. An adaptive network can rewire itself efficiently—that is, restructure or reorganize its topology on the basis of environmental shifts—to continue providing efficient performance. For example, if a supplier can't reliably meet a customer's demands, the customer should be able to choose another supplier.

A typical supply chain with a tree-like or hierarchical structure lacks these four properties—the clustering coefficient is nearly zero, and the characteristic path length scales linearly with the number of nodes (or agents) N. In designing complex agent networks with built-in survivability, conventional optimization tools won't work because of the problem's extremely large scale. When networks were smaller, we could understand their overall behavior by concentrating on the individual components' properties. But as networks expand, this becomes impossible, so we shift focus to the statistical properties of the collective behavior.

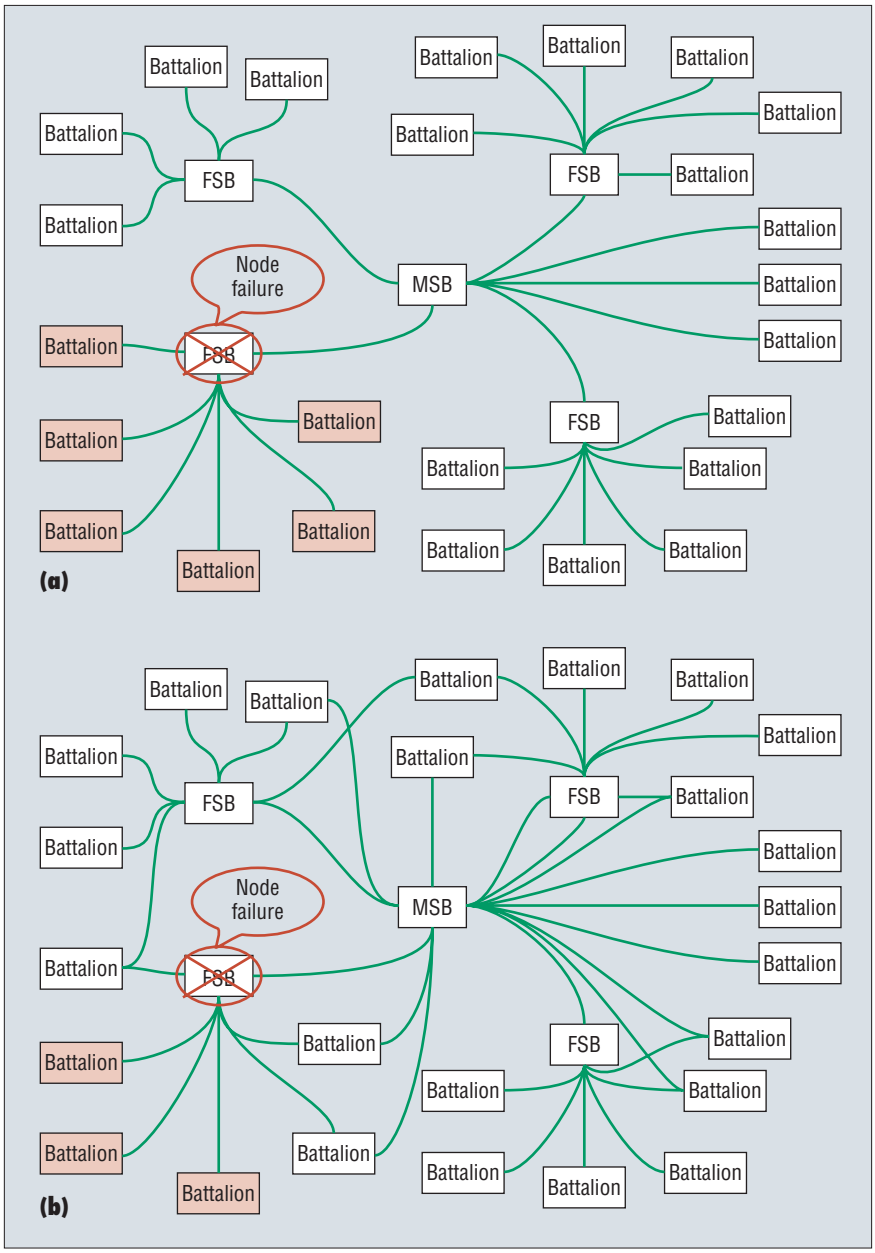

Figure 1. How redundancy affects survivability. (a) A part of the multiagent system for military logistics modeled using the UltraLog (www.ultralog.net) program. This example models each entity, such as main support battalion, forward support battalion, and battalion, as a software agent. (We've changed the agents' names for security reasons.) In the current scenario, MSBs send the supplies to the FSBs, who in turn forward these to battalions. (b) A modified military supply chain with some redundancy built into it. This network performs much better in the event of node failures and hence is more dependable than the first network.

Using topologies

Studying complex networks such as protein interaction networks, regulatory networks, social networks of acquaintances, and information networks such as the Web is illuminating the principles that make these networks extremely resilient to their respective chaotic environments. The core principles extracted from this exploration will prove valuable in building robust models for survivable complex agent networks.

Complex-network theory currently offers random-graph, small-world, and scale-free network topologies as likely candidates for survivable networks (see the sidebar "Complex Networks" for more on this topic). Evaluating these for survivability (see Figure 2), we find that no one topology consistently outperforms the others. For example, while small-world networks have better clustering properties, scale-free networks are significantly more robust to random attacks. So, we can't directly use these topologies to build supply networks. We can, however, use their evolution principles to build supply chain networks that perform well in all respects of the survivability components.

Researchers have studied complex networks in part to find ways to design evolutionary algorithms for modeling networks

Social scientists, among the first to study complex networks extensively, focused on *acquaintance networks*, where nodes represent people and edges represent the acquaintances between them. Social psychologist Stanley Milgram posited the "six degrees of separation" theory that in the US, a person's social network has an average acquaintance path length of six.[1] This turns out to be a particular instance of the small-world property found in many real-world networks, which, despite their large size, have a relatively short path between any two nodes.

An early effort to model complex networks introduced *random graphs* for modeling networks with no obvious pattern or structure.[2] A random graph consists of N nodes, and two nodes are connected with a connection probability p. Random graphs are statistically homogeneous because most nodes have a degree (that is, the number of edges incident on the node) close to the graph's average degree, and significantly small and large node degrees are exponentially rare.

However, studying the topologies of diverse large-scale networks found in nature reveals a more complex and unpredictable dynamic structure. Two measures quantifying network topology found to differ significantly in real networks are the *degree distribution* (the fraction of nodes with degree k) and the *clustering coefficient*. Later modeling efforts focused on trying to reproduce these properties.[3,4] Duncan Watts and Steven Strogatz introduced the concept of *small-world networks* to explain the high degree of transitivity (order) in complex networks.[5] The Watts-Strogatz model starts from a regular 1D ring lattice on L nodes, where each node is joined to its first K neighbors. Then, with probability p, each edge is rewired with one end remaining the same and the other end chosen uniformly at random, without allowing multiple edges (more than one edge joining a pair of vertices) or loops (edges joining a node to itself). The resulting network is a regular lattice when $p = 0$ and a random graph when $p = 1$, because all edges are rewired. This network class displays a high clustering coefficient for most values of p, but as $p \to 1$, it behaves like a random graph.

Albert-László Barabási and Réka Albert later proposed an evolutionary model based on growth and preferential attachment leading to a network class, *scale-free networks*, with power law distribution.[6] Many real-world networks' degree distribution follows a power law, fundamentally different from the peaked distribution observed in random graphs and small-world networks. Barabási and Albert argued that a static random graph of the Watts-Strogatz model fails to capture two important features of large-scale networks: their constant growth and the inherent selectivity in edge creation. Complex networks such as the Web, collaboration networks, or even biological networks are growing continuously with the creation of new Web pages, the birth of new individuals, and gene duplication and evolution. Moreover, unlike random networks where each node has the same chance of acquiring a new edge, new nodes entering the scale-free network don't connect uniformly to existing nodes but attach preferentially to higher-degree nodes. This reasoning led Barabási and Albert to define two mechanisms:

- *Growth:* Start with a small number of nodes—say, m_0—and assume that every time a node enters the system, m edges are pointing from it, where $m < m_0$.
- *Preferential attachment:* Every time a new node enters the system, each edge of the newly connected node preferentially attaches to a node i with degree k_i with the probability

$$\Pi_i = \frac{k_i}{\sum_j k_j}$$

Research has shown that the second mechanism leads to a network with power-law degree distribution $P(k) = k^{-\gamma}$ with exponent $\gamma = 3$. Barabási and Albert dubbed these networks "scale free" because they lack a characteristic degree and have a broad tail of degree distribution. Following the proposal of the first scale-free model, researchers have introduced many more refined models, leading to a well-developed theory of evolving networks.[7]

Protein-to-protein interactions in metabolic and regulatory networks and other biological networks also show a striking ability to survive under extreme conditions. Most of these networks' underlying properties resemble the three most familiar networks found in the literature (see Figure 1 in the article).

Complex networks are also vulnerable to node or edge losses, which disrupt the paths between nodes or increase their length and make communication between them harder. In severe cases, an initially connected network breaks down into isolated components that can no longer communicate. Numerical and analytical studies of complex networks indicate that a network's structure plays a major role in its response to node removal. For example, scale-free networks are more robust than random or small-world networks with respect to random node loss.[8] Large scale-free networks will tolerate the loss of many nodes yet maintain communication between those remaining. However, they're sensitive to removal of the most-connected nodes (by a targeted attack on critical nodes, for example), breaking down into isolated pieces after losing just a small percentage of these nodes.

References

1. S. Milgram, "The Small World Problem," *Psychology Today*, vol. 2, May 1967, pp. 60–67.
2. P. Erdös and A. Renyi, "On Random Graphs I," *Publicationes Mathematicae*, vol. 6, 1959, pp. 290–297.
3. S.N. Dorogovtsev and J.F.F. Mendes, "Evolution of Networks," *Advances in Physics*, vol. 51, no. 4, 2002, pp. 1079–1187.
4. M.E.J. Newman, "The Structure and Function of Complex Networks," *SIAM Rev.*, vol. 45, no. 2, 2003, pp. 167–256.
5. D.J. Watts and S.H. Strogatz, "Collective Dynamics of 'Small-World' Networks," *Nature*, vol. 393, June 1998, pp. 440–442.
6. A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, Oct. 1999, pp. 509–512.
7. R. Albert and A.-L. Barabási, "Statistical Mechanics of Complex Networks," *Reviews of Modern Physics*, Jan. 2002, pp. 47–97.
8. R. Albert, H. Jeong, and A.-L Barabási, "Error and Attack Tolerance of Complex Networks," *Nature*, July 2000, pp. 378–382.

with distinct properties found in nature. A network's evolutionary mechanism is designed such that the network's inherent properties emerge owing to the mechanism. For example, small-world networks were designed to explain the high clustering coefficient found in many real-world networks, while the "rich get richer" phenomenon used in the Barabási-Albert model explains the scale-free distribution.[2]

Similarly, we seek to design supply networks with inherent survivability components (see Figure 3), obtaining these components by coining appropriate growth mechanisms. Of course, having all the aforementioned properties in a network might not be practically feasible—we'd likely have to negotiate trade-offs depending on the domain. Also, domain specificities might make it inefficient to incorporate all properties. For instance, in a supply network, we might not be able to rewire the edges as easily as we can in an information network, so we would concentrate more on obtaining other properties such as low characteristic path length, robustness to failures and attacks, and high clustering coefficients. So, the construction of these networks is domain specific.

Establishing edges between network nodes is also domain specific. For instance, in a supply network, a retailer would likely prefer to have contact with other geographically convenient nodes (distributors, warehouses, and other retailers). At the same time, nodes in a file-sharing network would prefer to attach to other nodes known to locate or hold many shared files (that is, nodes of high degree).

Obtaining the survivability components

While evolving the network on the basis of domain constraints, we need to incorporate four traits into the growth model for obtaining good survivability components.

The first is *low characteristic path length*. During network construction, establish a few long-range connections between nodes that require many steps to reach one from another.

The second is *good clustering*. When two nodes A and B are connected, new edges from A should prefer to attach to neighbors of B, and vice versa.

The third is *robustness to random and targeted failure*. Preferential attachment—where new nodes entering the network don't connect uniformly to existing nodes but attach preferentially to higher-degree nodes (see the side-

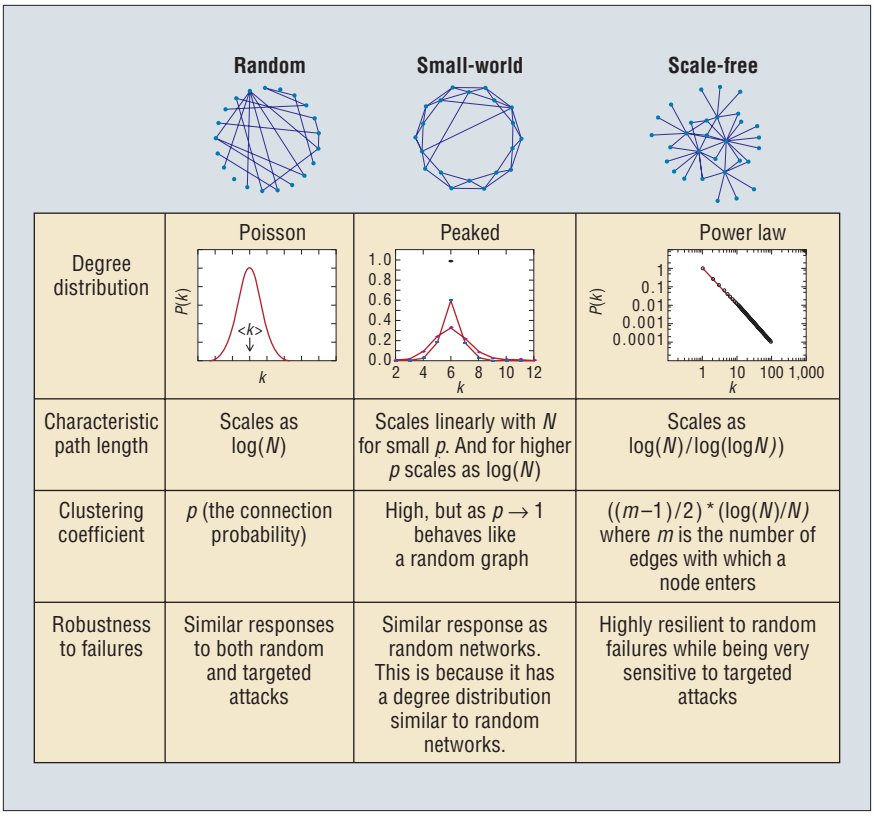

Figure 2. Comparing the survivability components of random, small-world, and scale-free networks.

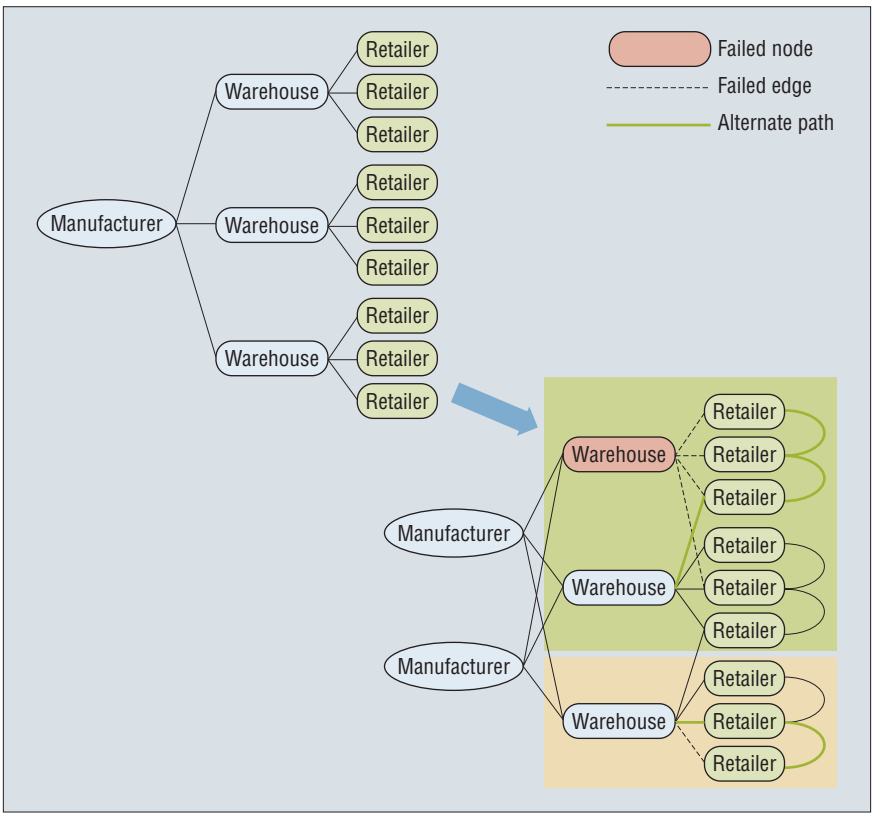

Figure 3. The transition from supply chain to a survivable supply network.

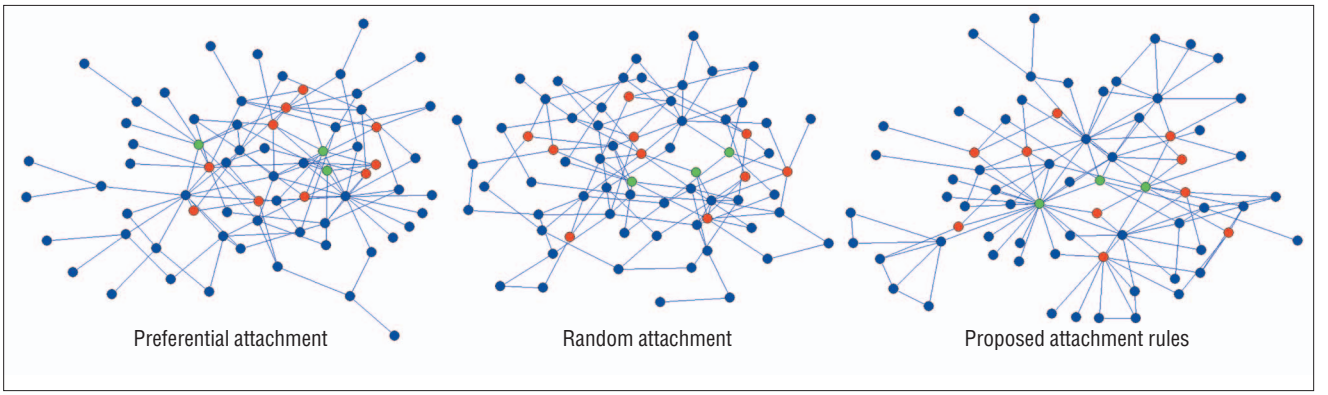

Figure 4. Snapshots of the modeled networks during their growth, where the nodes number 70. MSBs are green, FSBs are red, and battalions are blue.

bar for more details)—leads to scale-free networks with very few critical and many not-so-critical nodes. Here we measure a node's criticality in terms of the number of edges incident on it. So, these networks are robust to random failures (the probability that a critical node fails is very small) but not to targeted attacks (attacking the very few critical nodes would devastate the network). Also, it's not practically feasible to have all nodes play an equal role in the system—that is, be equally critical. Thus, the network should have a good balance of critical, not-so-critical, and noncritical nodes.

The fourth is *efficient rewiring*. Rewiring edges in a network might or might not be feasible, depending on the domain. But where it is feasible, it should preserve the other three traits.

Although complete graphs come equipped with good survivability components, they clearly aren't cost effective. Allowing every agent in an agent system to communicate with every other agent uses system bandwidth inefficiently and could completely bog down the system. So the amount of redundancy results from a trade-off between cost and survivability.

An illustration

Suppose we want to build a topology for a military supply chain that must be survivable in wartime. First, we broadly classify the network nodes into three types:

- *Battalions* prefer to attach to a highly connected node so that the supplies from different parts of the network will be transported to them in fewer steps. Battalions also require quick responses, so they prefer the subsequent links to attach to nodes at convenient shorter distances (in our model we considered a fixed distance of two).
- A *forward support battalion* prefers to attach to highly connected nodes so that its supplies proliferate faster in the network. The supply range from an FSB goes up to a particular distance (at most three in our model).
- A *main support battalion* also prefers to attach to a highly connected node to enable its supplies to proliferate faster in the network. We assume an unrestricted supply reach from an MSB, thus facilitating some long-range connections.

In a conventional logistics network, the MSBs supply commodities (such as ammunitions, food, and fuel) to the FSBs, who in turn forward them to the battalions. Our approach doesn't restrict node functionalities as such—for example, we assume that even a battalion can supply commodities to other battalions if necessary.

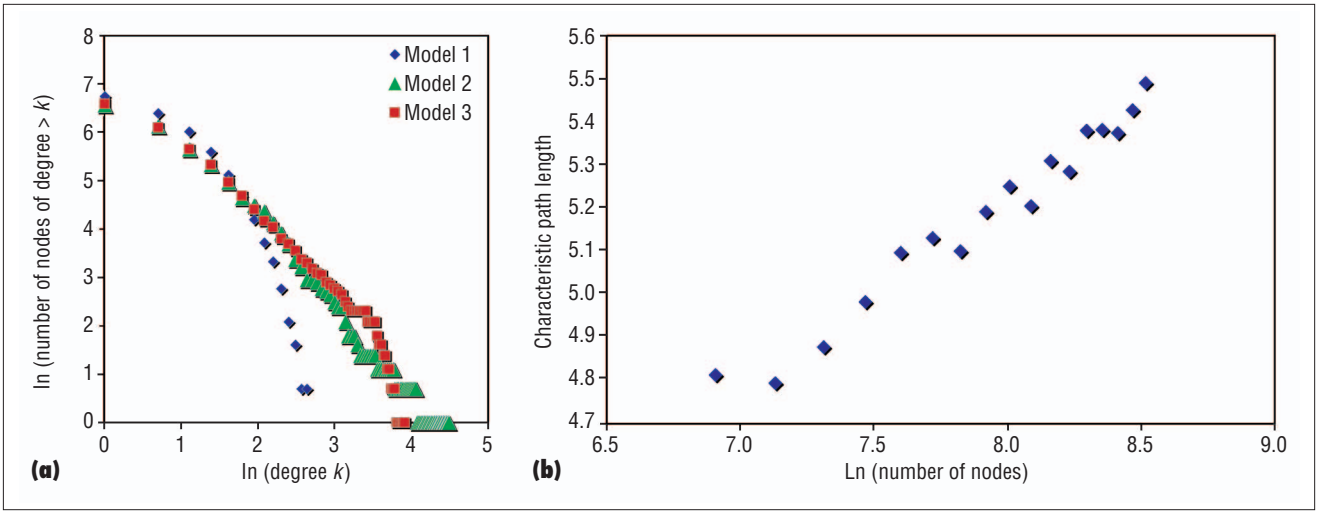

Figure 5. How our proposed network performed: (a) the log-log of the degree distribution for all the three networks; (b) the characteristic path length of the proposed network against the log of the number of nodes.

Table 1. Simulation results.

	Model 1 (random)	Model 2 (preferential)	Model 3 (proposed)
Clustering coefficient	0.0038–0.0039	0.013–0.019	0.35–0.39
Characteristic path length	5.26–5.36	4.09–4.25	4.69–4.79

Growth mechanisms

Start with a small number of nodes—say, m_0—and assume that every time a node enters the system, m edges are pointing from it, where $m < m_0$. Battalions, FSBs, and MSBs enter the system in a certain ratio $l{:}m{:}n$ where $l > m > n$:

- A battalion has one edge pointing from it and a second edge added with a probability p.
- An FSB has three edges pointing from it.
- An MSB has five edges pointing from it.

The attachment rules applied depend on which node type enters the system:

- For a battalion, the first edge attaches to a node i of degree k_i with the probability

$$\Pi_i = \frac{k_i}{\sum_j k_j}.$$

  The second edge, which exists with a probability p, attaches to a randomly chosen node at a distance of two.
- For an FSB, the first edge attaches to a node i of degree k_i with the probability

$$\Pi_i = \frac{k_i}{\sum_j k_j}.$$

  The subsequent edges attach to a randomly chosen node at a distance of at most three.
- For an MSB, each edge attaches preferentially to a node i with degree k_i with the probability

$$\Pi_i = \frac{k_i}{\sum_j k_j}.$$

Simulation and analysis

Using this method, we built a network of 1,000 nodes with l, m, and n being 25, 4, and 1 (we obtained these values from the current configuration of the military logistics network used in the UltraLog program) and $p = 1/2$. We compared this network's survivability with that of two other networks built using similar mechanisms except that one used purely preferential attachment rules (similar to scale-free networks) and the other used purely random attachment rules (similar to random networks) (see Figure 4). All three networks had an equal number of edges and nodes to ensure fair comparison.

We refer to the networks built from random, preferential, and proposed attachment rules as Models 1, 2, and 3, respectively. As we noted earlier, a typical military supply chain (see Figure 1a) with a tree-like or hierarchical structure has deficient survivability components, making it vulnerable to both random and targeted attacks. Models 1, 2, and 3 outperform the typical supply network in all survivability components.

Figure 5a shows the three models' degree distribution. As expected, the preferential-attachment network has a heavier tail than the other two networks. We measured survivability components for all three networks.

The clustering coefficient for Model 3 was the highest (see Table 1). The Model 3 attachment rules, especially those for battalions and FSBs, contribute implicitly to the clustering coefficient, unlike the attachment rules in the other models.

The proposed network model's characteristic path length measured between 4.69 and 4.79 despite the network's large size (1,000 nodes). This value puts it between the preferential and random attachment models. Also, as Figure 5b shows, the characteristic path length increases in the order of $\log(N)$ as N increases. Model 3 clearly displays small-world behavior.

To measure network robustness, we removed a set of nodes from the network and evaluated its resilience to disruptions. We considered two attacks types: random and targeted. To simulate random attacks, we removed a set of randomly chosen nodes; for targeted attacks, we removed a set of nodes selected strictly in order of decreasing node degree. To determine robustness, we measured how the size of each network's largest connected component, characteristic path length, and maximum distance within the largest connected component changed as a function of the number of nodes removed. We expect that in a robust network the size of the largest connected component is a considerable fraction of N (usually $O(N)$), and the distances between nodes in the largest connected component don't increase considerably.

For random failures, Figure 6 shows that Model 3's robustness nearly matches that of the preferential-attachment network (note that scale-free networks are highly resilient to ran-

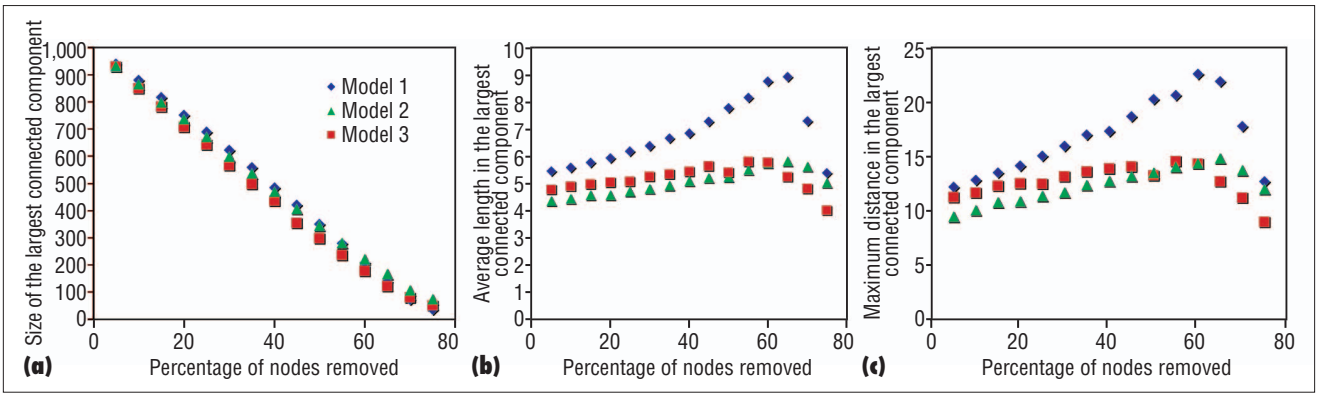

Figure 6. Responses of the three networks to random attacks, plotted as (a) the size of the largest connected component, (b) characteristic path length, and (c) maximum distance in the largest connected component against the percentage of nodes removed from each network.

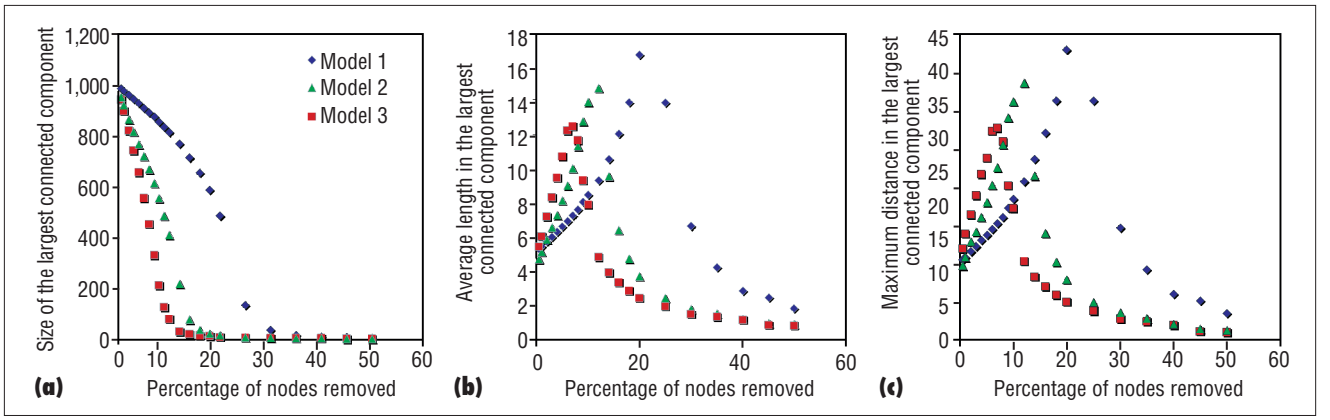

Figure 7. The three networks' responses to targeted attacks, plotted as (a) the size of the largest connected component, (b) characteristic path length, and (c) maximum distance in the largest connected component against the percentage of nodes removed from each network.

dom failures). Also, the decrease in the largest connected component's size is linear with respect to the number of nodes removed, which corresponds to the slowest possible decrease. So, we can safely conclude that these networks are robust to random failures—most of the nodes in the network have a degree less than four, and removing smaller-degree nodes impacts the networks much less than removing high-degree nodes (called *hubs*).

These networks' responses to targeted attacks are inferior compared to their resilience to random attacks (see Figure 7). The size of the largest component decreases much faster for the proposed network than for the other two networks, but the proposed network performs better on the other two robustness measures. That is, the distances in the connected component are considerably smaller when more than 10 percent of nodes are removed.

We can improve robustness to targeted attacks by introducing constraints in the attachment rules. Here we assume that node type constrains its degree—that is, network MSBs, FSBs, and battalions can't have more than m_1, m_2, and m_3 edges, respectively, incident on them. This is a reasonable assumption because in military logistics (or any orga-

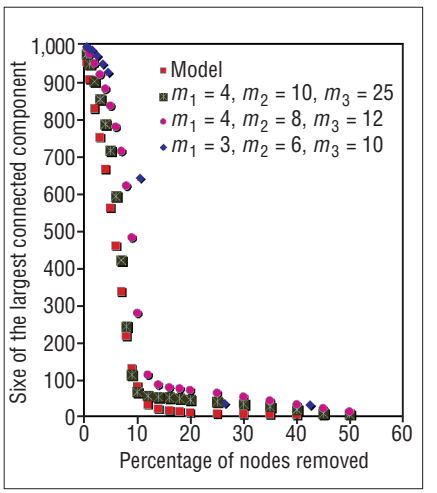

Figure 8. The proposed network's responses to targeted attacks for different values of m_1, m_2, and m_3.

The Authors

Hari Prasad Thadakamalla is a PhD student in the Department of Industrial and Manufacturing Engineering at Pennsylvania State University, University Park. His research interests include supply networks, search in complex networks, stochastic systems, and control of multiagent systems. He obtained his MS in industrial engineering from Penn State. Contact him at hpt102@psu.edu.

Usha Nandini Raghavan is a PhD student in industrial and manufacturing engineering at Pennsylvania State University, University Park. Her research interests include supply chain management, graph theory, complex adaptive systems, and complex networks. She obtained her MSc in mathematics from the Indian Institute of Technology, Madras. Contact her at uxr102@psu.edu.

Soundar Kumara is a Distinguished Professor of industrial and manufacturing engineering. He holds joint appointments with the Department of Computer Science and Engineering and School of Information Sciences and Technology at Pennsylvania State University. His research interests include complexity in logistics and manufacturing, software agents, neural networks, and chaos theory as applied to manufacturing process monitoring and diagnosis. He's an elected active member of the International Institute of Production Research. Contact him at skumara@psu.edu.

Réka Albert is an assistant professor of physics at Pennsylvania State University and is affiliated with the Huck Institutes of the Life Sciences. Her main research interest is modeling the organization and dynamics of complex networks. She received her PhD in physics from the University of Notre Dame. She is a member of the American Physical Society and the Society for Mathematical Biology. Contact her at ralbert@phys.psu.edu.

Table 2. The proposed network's characteristic path length for different m_1, m_2, and m_3 values.

Values of m_1, m_2, and m_3	Characteristic path length
$m_1 = \infty$, $m_2 = \infty$, $m_3 = \infty$	4.4
$m_1 = 4$, $m_2 = 10$, $m_3 = 25$	6.2
$m_1 = 4$, $m_2 = 8$, $m_3 = 12$	7.1
$m_1 = 3$, $m_2 = 6$, $m_3 = 10$	8.0

nization's logistics management, for that matter), the suppliers might not be able to cater to more than a certain number of battalions or other suppliers. Initial experiments (see Figure 8) show that a network with these constraints displayed improved robustness to targeted attacks while not deviating much from the clustering coefficient. However, as we restrict how many links a node can receive, the network's characteristic path length increases (see Table 2). Clearly a trade-off exists between robustness to targeted attacks and the average characteristic path length.

The fourth measure of survivability, network adaptivity, relates more to node functionality than to topology. Node functionality should facilitate the ability to rewire. For example, if a supplier can't fulfill a customer's demands, the customer seeks an alternate supplier—that is, the edge connected to the supplier is rewired to be incident on another supplier. Our model rewires according to its attachment rules. We conjecture that in such a case, other survivability components (clustering coefficient, characteristic path length, and robustness) will be intact. But to make a stronger argument we need more analysis in this direction.

The growth mechanism we describe is more like an illustration because real-world data aren't available, but we can always modify it to incorporate domain constraints. For example, we've assumed that a new node can attach preferentially to any node in the network, which might not be a realistic assumption. If specific geographical constraints are known, we can modify our mechanism to make the new node entering the system attach preferentially only within a set of nodes that satisfy the constraints. ■

Acknowledgments

We thank the anonymous reviewers for their helpful comments. We acknowledge DARPA for funding this work under grant MDA972-01-1-0038 as part of the UltraLog program.

References

1. J.M. Swaminathan, S.F. Smith, and N.M. Sadeh, "Modeling Supply Chain Dynamics: A Multiagent Approach," *Decision Sciences*, vol. 29, no. 3, 1998, pp. 607–632.
2. A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, Oct. 1999, pp. 509–512.