Towards an Adaptive Economic Society of Peers

Paul E. Silvey
The MITRE Corporation
psilvey@mitre.org

Abstract

This paper presents our on-going research into self-organizing and self-healing networks of resource brokers, for use in complex Peer-to-Peer (P2P) resource sharing and Grid computing application environments. It discusses our objectives for an efficient, robust, and decentrally controlled infrastructure, and describes our efforts to understand the co-dependencies between protocols, topologies, and continuously changing environmental conditions. Motivated by dynamic social and economic systems, our research project is nicknamed AESOP, for Adaptive Economic Society Of Peers. We discuss our specific investigations into adaptive topologies for flooding-based distributed search protocols, Genetic Algorithm explorations seeking good topological structures, and current work to develop an agent-based simulation to investigate the effects of peer transience on P2P join and topology maintenance protocols.

1. Introduction

Many researchers of complex systems have taken an interest in the general nature of dynamic networks, with the realization that discovered principles may have great influence on a broad range of other scientific investigations. For computer scientists and software engineers, P2P and Grid computing architectures are emerging as a viable technical approach to the construction of massively distributed information processing and file sharing systems [8]. These distributed resource discovery and sharing applications depend on large networks of cooperative and autonomous computers to provide their infrastructures. Overall performance of these systems depends on the interconnection structure or overlay topology of the network graph, as well as the communication protocols used to perform distributed coordination and work. Topologies that are well suited to these kinds of applications are not well understood, due to fundamental combinatorial complexities and the multitude of potentially relevant dynamic variables. For example, the number of possible ways to connect the nodes of an undirected graph of a given size is prohibitively large to be investigated directly or easily studied. While we know a great deal about certain topological families from graph theory and empirical analysis of existing real-world networks, much remains to be learned.

One remarkable property of many naturally occurring networks is that they conform to a topology known as *scale-free*, where the degree distribution (number of links per node) follows a power-law function. Most nodes in the network are sparsely connected, but a very few are extremely well connected, acting as important hubs for keeping the network together and for keeping the path distance between peers bounded. This kind of a topology is efficient for P2P message propagation, but efficiency comes at the expense of robustness, which is an important requirement for many proposed P2P applications. When a few of the largest hubs cease to function, due to random failures or the malicious behavior of adversaries, the whole network can fail catastrophically. Further, it appears that scale-free topologies are so easy to construct that they can happen as accidental side effects of distributed network connection protocols, and that special effort must be exerted to prevent them in many cases. Yet balanced degree (regular graph) network topologies tend to require greater bandwidth and have larger diameters (longest direct path between any pair of nodes in the graph). Large diameter graphs imply longer query response times, or artificial horizon effects from message time-to-live propagation strategies.

As others have observed [2], random networks, though frequently studied or assumed, do not often occur naturally, and scale-free networks are not ideal if robustness is important. Our research is therefore seeking qualitative insights about alternate network topologies that are good for P2P applications, with the hope that we can actively manage a dynamic network of peer connections without imposing too rigid a structure or incurring too much overhead in the process.

Although P2P is most often associated with file sharing applications, and Grid computing with distributed management of computer system resources, they have much in common. Both P2P and Grid systems benefit from the inherent parallelism of the network, and they both require mechanisms for distributed discovery of available resources as a precursor to their effective use. We use the term *resource* to include information objects such as data files or multimedia streams, as well as computer disk storage, available CPU cycles, sensors, or even web services. Thus the focus of our research is on the core problems associated with distributed resource brokering in complex network environments. We are interested in removing the vulnerable single points of failure that centralized control strategies depend on, or that scale-free networks present through their very topological structure. Consequently, our research is first of all aimed at understanding what makes a good P2P network topology, and then devising topology monitoring and adaptation protocols that will enable the network to be self-organizing and self-healing in the face of environmental dynamics. One important type of environmental change is known as *peer transience*, which is characterized by the rate of peer arrivals and departures. Other dynamic forces for P2P systems include economic fluctuations in resource supply, demand, and capacity. Taken together, these objectives suggested the project nickname, as we seek to develop robust, scalable, and efficient mechanisms for an Adaptive Economic Society Of Peers (AESOP).

This paper is structured as follows. Section 2 discusses the background of our research and other related work in distributed discovery. In section 3 we discuss the use of a Genetic Algorithm (GA) to explore topological space, and lessons learned from that exercise. Section 4 concludes with some on-going work to develop an agent-based modeling and simulation test-bed, which we are using to empirically evaluate join protocols and to study the effects of peer transience. Finally, we list our references in section 5.

2. Background and Related Research

MITRE has been working on global information management problems for a number of years, based on flexible and distributed metadata management, to support autonomous and loosely-coupled resource providers, intermediaries, and consumers [5]. While centralized brokering strategies present a number of hard problems in their own right, scalability and robustness requirements led us to pursue distributed search solutions based on P2P concepts. For our purposes, we are interested in networks that are robust against attacks and/or component failures, while simultaneously being efficient in time and space performance characteristics. We study undirected graphs to ensure that pair-wise connections between nodes have been mutually established, perhaps based on authentication, access-control, or trust relationships.

The distributed resource discovery problem has been addressed several different ways, depending on the nature of the items sought and the structure and stability of the network. One simple kind of search problem is known as *key-indexed lookup*, where each resource has a uniquely assigned identifier. The distributed version of this uses a network of computers to provide collective and scalable storage. A number of proposed solutions are based on Distributed Hash Tables (DHT), which have become an active area of research [3][7]. In addition to assuming unique resource identity, these approaches often impose a good deal of structure on the network topology, and assume that node transience is not too great.

A variation on distributed key-index lookup has also been investigated by social network researchers, where the problem is the classic one studied originally by Milgram of routing a physical message to a particular person based on limited information about them [13]. In this case, there is a unique copy of the “target”, and the search is conducted by sequential token passing, guided from node to neighbor node using local best guesses of how to get progressively closer. Here, routing decisions are guided by heuristics, versus the more deterministic addressing mechanisms of DHTs. Similar serialization of the search process is also characteristic of various P2P strategies known as *walkers*, including simple random walk searching. Variations on this approach include allowing more than one simultaneous walker thread [12], and adding preferential biases toward neighbor nodes with high degree [1] or high capacity [4]. More general than simple social network search, walker strategies can of course be used when there is more than one possible target, or when matching is partial or inexact. Nevertheless, these strategies often assume that the topology has certain fixed characteristics such as being random or scale-free, and are primarily aimed at reducing communication loads on the network, albeit at the expense of a more timely response.

When looking for an information resource such as a particular multimedia file of which there may be many copies distributed throughout the network, we would like to find at least one of them as quickly as possible. In fact, since replicas can be reliably detected using something like a hash code signature, it can be beneficial to find multiple copies, to enable a “swarming” (piecemeal parallel) download, saving time and distributing the server and bandwidth burden over multiple machines. Furthermore, if the resource being sought is only approximately described by the query (as is typical in information retrieval applications, say using keywords), we would also like to reach all nodes, but for a different reason – so that the nominated “hits” can be ranked and evaluated by the node initiating the search. This is the type of search problem for which parallel flooding-based search strategies such as Gnutella were designed [9].

Query messages propagate to nodes in the network as fast as the topology allows, bounded by the diameter of the graph. In practice, however, the unstructured and rather uncontrolled nature of real topologies results in excessive bandwidth use, and has prompted developers of flooding-based distributed search systems to impose time-to-live constraints on messages to limit the cost to the network. More recently, additional techniques that differentiate client edge peers from internal server peers (called *super peers* or *ultra peers*) have also been used to help with the bandwidth demand problem, by restricting flooding to the small subset of these interconnected internal nodes with greater capacities and longevity.

Our objectives have been to support fast and thorough discovery based on partial and inexact metadata matching, in dynamic environments with the potential for high node transience. As a result, our research has been focused on improving the performance of flooding-based distributed search through a strategy of active topology adaptation. Other researchers have pursued similar objectives by controlling the join process [14], as well as by adapting existing topologies [4]. In our initial AESOP experiments, we conducted simulations of adaptive strategies to evaluate two separate heuristics, known as *edge thinning* and *diameter folding* [15]. We were able to demonstrate how these strategies provide tangible system performance improvements for a variety of starting topologies, but scale-free networks proved challenging to avoid. Based on these investigations, we observed that the best performing topologies had three important characteristics. First, they were sparse in their connections, but not as minimal as a spanning tree. Second, they were balanced, being regular or nearly regular in degree, with each node having roughly the same small number of links[1]. And third, they were tight, with small diameters. They also tended to appear approximately symmetric, but this feature was not formally evaluated.

Note however, that these are simply descriptive topological properties, involving measurable attributes for any given graph. Beyond heuristics such as those we have already employed, the desirable graph properties do not by themselves suggest how we might change a topology to make it better. Furthermore, as we discovered, our original adaptive strategies had particular problems with common scale-free networks. As a result, we decided to try an experiment using a Genetic Algorithm [10][11], to explore the vast search space of possible topologies, using a fitness function based on the goodness properties listed above. The GA performs a parallel search function based on evolutionary principles, and we used it with the hope that it might uncover topological properties that our intuitions alone would not suggest. It is therefore simply a tool to ultimately assist us in protocol design, and not itself a candidate for P2P resource discovery. For further details on how GAs in general work, the reader is referred to the literature. The specifics of our experiments are outlined next.

3. Genetic Algorithm Suggests Girth

To formulate our problem for genetic algorithm exploration, we needed to represent each individual topology as a bit-string chromosome, and to define a formal fitness measure. We used a bit-string form of the graph's adjacency matrix as our chromosome, with a 1 bit in the (i, j) position if a link exists in the graph between node i and node j, and a 0 bit otherwise. An edge map for an undirected graph is a symmetric bit matrix, so to save space and allow all possible bit patterns, we used only the string of bits from the upper triangular part of the adjacency matrix, which is $n(n\text{-}1)/2$ bits long for a network of n nodes.

Our fitness score was based on some preliminary analysis of random graphs, which showed that a nearly equal weighting between diameter and mean degree would work, since we desired to minimize them both simultaneously. Initially, we tried simply adding the graph diameter value to the mean degree value and sorting these fitness values in ascending order. Note that smaller values of this metric therefore represent greater fitness in the population.

$$\mathrm{F}_{\mathrm{orig}} = \mathrm{diam}(G) + \mathrm{deg}_{\mu}(G)$$

However, after testing for a few thousand generations, we noted that all members of the population had the same identical score, so our fitness measure was not an adequate differentiator for selection. As a result, we added the node degree standard deviation value to the fitness score, since a zero standard deviation for degree occurs exactly when the graph has a regular topology, and is small for nearly balanced graphs. Taken together, this fitness score codes for time efficiency (low diameter), space efficiency (low total number of edges as represented by mean degree), and robustness (regular or nearly balanced topologies).

$$\mathrm{F}_{\mathrm{new}} = \mathrm{diam}(G) + \mathrm{deg}_{\mu}(G) + \mathrm{deg}_{\sigma}(G)$$

Using the fitness score just described, we ran GA experiments using graphs of size 16, 32, 64, and 128 nodes. For each size graph, we maintained a population of 100 individuals, and ran the GA for at least 10,000 generations. The 128 node graphs continued to show fitness score improvements after the first 10,000

---

[1] Completely balanced or regular graphs are not likely to occur in real P2P networks, where wide ranges in capabilities and peer transience exist. They provide useful idealized models, but constraining to a *range* of degrees is more feasible in practice.

generations, so we continued them for 23,000 generations total. On each generation, the best scoring 10% of the population was retained for the next generation, and another 10% were newly created random graphs. The remaining 80% were generated using single point crossover and mutation operators. Pairs of new individuals were generated from randomly selected pairs of parents, where the parent selection was biased toward the best scoring individuals from the current generation. This bias was accomplished by sorting the population in ascending fitness order, and selecting a parent whose index was the scaled value obtained by multiplying two uniform random numbers in the interval [0, 1).

We experimented with a few different values for mutation rate, but ended up using 2% for most of the runs. We allowed multiple mutations to occur to each offspring generated, continuing as long as the appropriately weighted coin-flip signaled that a mutation was to be made. Finally, we only kept offspring that represented connected graphs, generating candidates until the population quota for the next generation had been filled. At various generation milestones and at the end of our runs, we externalized the graphs in a form suitable for subsequent analysis. In addition, various statistics were collected and visualizations produced.

In addition to running GA evolutions on random initial populations of the four different sized graphs, we also ran a few thousand generations on initial populations generated to be random near-regular topologies whose degree target was the rounded natural log of the number of nodes. These topologies were expected to do well based on previous analysis, and the fact that they almost always were regular and therefore had a zero value for the degree standard deviation.

Visual inspection of the resulting graphs from the smallest test cases suggested that good topologies had another property that we had not seen before. Each edge participated in a cycle, and these cycles tended to be fairly large. On reflection, this property made sense, since messages flooded from a given node along multiple paths would not soon reunite at a common node if the cycles were large, and the network would therefore more efficiently propagate queries. In graph theory, the size of the minimal cycle in a graph is called its *girth*. So it appeared that graphs with high girth might be desirable for P2P networks as well. Following this observation, we discovered the class of graphs known as *cubic symmetric graphs* that have been sought after and collected by graph theorists in what is known as the Foster Census [6]. A 128-node graph from this collection is depicted in Figure 1. It is called *cubic* because it has regular degree of 3, and this particular visual layout makes it easy to see its symmetrical structure. This graph, whose Foster Census ID is F128B, has a girth of 10 and a diameter of 8, and is the most robust and efficient graph of this size we have seen for flooding-based distributed search.

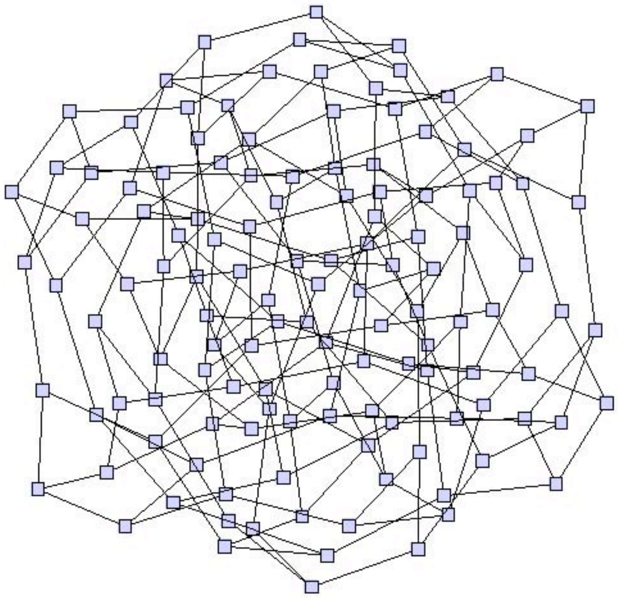

Figure 1 – Cubic Symmetric Graph F128B

The property of high girth, like the other good topology characteristics, did not by itself offer a prescription for topology adaptation, but did suggest a new edge thinning heuristic. Our initial method for edge thinning was based on message passing utility scores, aimed at locating edges that offered excessive redundancy or caused messages to arrive late because they were part of odd-length cycles [15]. The girth insight suggested that we might instead nominate edges for removal when they were part of short cycles.

Table 1 – Topology Adaptation Statistics
for Short Cycle Thinning + Diameter Folding
(initial values followed by shaded results)

Initial Graph Topology	Degree			Path			Search		
	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
Circle Net	2	2.0	2	1	32.3	64	128	128.0	128
	2	3.0	4	1	4.8	8	218	226.7	237
Binary Hypercube	7	7.0	7	1	3.5	7	448	448.0	448
	3	3.7	4	1	3.9	7	235	235.0	235
Random Regular: degree 3	3	3.0	3	1	5.2	9	215	227.0	236
	3	3.0	3	1	4.8	8	219	226.8	237
Random Regular: degree 7	7	7.0	7	1	2.7	4	615	639.3	662
	3	3.5	4	1	4.0	7	248	262.9	280
Random Net	2	7.4	13	1	2.6	5	634	665.7	695
	3	3.4	4	1	4.1	7	250	264.0	277
Scale-Free Net (3 link join rule)	3	5.9	20	1	2.9	5	494	515.1	542
	3	3.5	4	1	4.1	7	249	261.4	284
Scale-Free Net (6 link join rule)	6	11.7	29	1	2.3	4	1061	1108.7	1166
	3	3.4	4	1	4.1	7	244	259.6	278
Full Net	127	127.0	127	1	1.0	1	16129	16129.0	16129
	3	3.5	4	1	4.0	7	247	265.6	280

Using this new thinning heuristic, we repeated our previous variable link-density experiment with the same 8 starting topologies and 128-node graph size. The adaptation simulations were run for thousands of iterations and converged in all cases as before. When more than one edge was an equally good candidate for removal or insertion, our previously developed *balancing bias* was again used (preference for link changes that minimize the degree standard deviation) [15]. The results of this repeated experiment confirmed that nominating short-cycle edges for replacement was superior for topologies that had nodes with high degrees, including scale-free networks, while being equally good in other respects. Contrary to our original experiment where scale-free and full networks retained some relatively high degree hubs, all starting topologies in the subsequent experiment were able to consistently adapt to near-regular degrees with tight upper bounds (see Table 1). As a result, short cycle thinning appears able to increase the efficiency and robustness of a broader range of possible starting topologies.

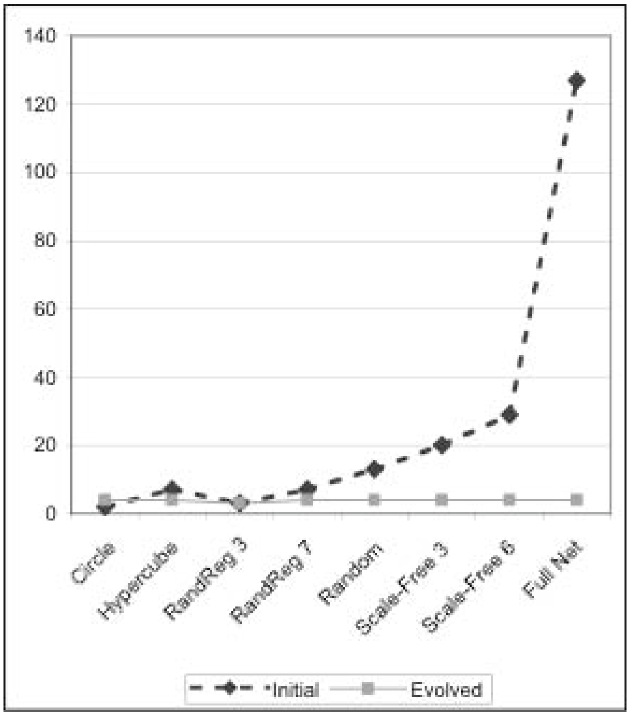

Figure 2 – Maximum Degree
for Short Cycle Thinning + Diameter Folding

4. Agent-Based Modeling and Simulation

The next stage of our research plan is to add more complexity to our simulations. In particular, we are developing an agent-based modeling and simulation test-bed that will let us repeat our adaptive topology experiments using a discrete event simulation framework. Our previous studies were based on a graph-centric simulation, and took advantage of global states of knowledge that individual peers would not automatically have in a real P2P system. In our new simulation environment for example, graph diameters are only ever knowable as estimates whose validity fades with time, especially in highly transient situations. These estimates must be inferred or communicated, and individual peers must make local decisions about topology adaptation operations they wish to perform. In order to support these additional complexities, we are designing some new protocols specifically to support topology adaptation and maintenance, while working to efficiently share useful topology state information as much as possible through the existing distributed search protocol.

For example, we have developed a variation on the flooding query that peers can occasionally use when initiating a query, which will let them establish the length of cycles that the edges between themselves and their neighbors participate in. As a result, peers will be able to gather this potentially important information as a side-effect of doing the ordinary work of discovery, at a slight cost in bandwidth and timeliness. Our simulation studies will be aimed at quantifying the specific costs vs. benefits of employing these adaptive topology information gathering and sharing protocols.

In addition, the new experimental test-bed will allow us to finally evaluate the effects of peer transience on our topology adaptation strategies. As part of this, we are implementing a particular published P2P join protocol that has been designed to create nearly balanced, connected graphs with small diameter [14]. This protocol currently relies on a centralized cache node, and we would like to explore ways to make it more de-centralized, and to empirically evaluate its theoretical claims in various simulated real world settings.

Future work is also planned for adding economic differentiators and value-driven peer decision-making. This will include modeling communities of interest, resource supply and demand fluctuations, and transience by peer type (e.g., always-on server-like peers vs. volatile mobile peers). We have particular long-term interest in understanding how to better enable global economic benefits from potentially selfish local social interactions among peers.

We have been encouraged to find that a P2P network topology can adapt to a form that makes flooding-based query routing tolerable, since this method of search is more thorough and potentially faster than selective walker-like strategies. Furthermore, the topologies we are producing via adaptive re-wiring of the network can actually benefit most query and advertisement routing strategies because they produce graphs that are sparse, balanced, and small in diameter.

Currently, many different technical approaches are independently being pursued for designing distributed dynamic systems with loose couplings. We envision a day, however, when many of these techniques will be effectively combined in hybrid forms. If P2P systems were able to make run-time tradeoffs based on immediate timeliness or bandwidth consumption considerations, they might choose a walker-based distributed search strategy in certain cases and flooding in others. Likewise, shaping the P2P topology to suit the environment can either be done at design time, for relatively static and well-understood environments, or be allowed to adapt dynamically when the expected range of behavior so warrants. Our research is directed at understanding the broad set of complex interactions that can occur between dynamic and autonomous nodes in large-scale P2P and Grid resource management systems.

5. References

[1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman, “Search in Power-Law Networks”, Physical Review E 64, 046135, 2001.

[2] R. Albert and A. L. Barabási, Statistical Mechanics of Complex Networks, Reviews of Modern Physics, 74 pp.47-97, 2002.

[3] M. Castro, P. Druschel, Y. C. Hu, A. I. T. Rowstron, “Topology-Aware Routing in Structured Peer-to-Peer Overlay Networks”, Future Directions in Distributed Computing, 2003, pp.103-107.

[4] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker, “Making Gnutella-like P2P Systems Scalable”, In Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany, August 2003.

[5] R. D. Cherinka and R. W. Miller, “Using XML-based Web Services to Enhance Information Interoperability in the Department of Defense”, Presented at the XML Conference & Exposition, Walt Disney World Dolphin Hotel, Orlando Florida, December 9-14, 2001.

[6] Cubic symmetric graphs (The Foster Census) http://www.cs.uwa.edu.au/~gordon/remote/foster/

[7] S. El-Ansary, L. O. Alima, P. Brand, S. Haridi, “Efficient Broadcast in Structured P2P Networks”, 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03), Berkeley, CA, February 2003.

[8] I. Foster, C. Kesselman, and S. Tuecke. “The anatomy of the grid: Enabling scalable virtual organizations”, International Journal of Supercomputer Applications, 15(3), 2001.

[9] The Gnutella Protocol Specification, version 0.4, rev 1.2. http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf

[10] J. H. Holland, Adaptation in Natural and Artificial Systems. Ann Arbor, MI: The University of Michigan Press, 1975.

[11] A. Kumar, R. M. Pathak, and M. C. Gupta, “Genetic algorithm based approach for designing computer network topology”, In Proceedings of the 1993 ACM Conference on Computer Science, Indianapolis, IN, USA, pp.358-365, 1993.

[12] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, “Search and Replication in Unstructured Peer-to-Peer Networks”, In Proceedings of 16th ACM International Conference on Supercomputing (ICS'02), New York, USA, June 2002.

[13] S. Milgram. The small world problem. Psychology Today, 2:60-67, 1967.

[14] G. Pandurangan, P. Raghavan, and E. Upfal, “Building Low-Diameter P2P Networks”, IEEE Journal on Selected Areas in Communication, Volume 21, Issue 6, August 2003.

[15] P. Silvey and L. Hurwitz, “Adapting Peer-to-Peer Topologies To Improve System Performance”, In Proceedings of the Hawai’i International Conference on System Sciences, Big Island, Hawaii, January 5–8, 2004.