Towards Autonomic Control of Network Topologies

Vincent Matossian and Manish Parashar

The Applied Software Systems Laboratory
Department of Electrical and Computer Engineering
Rutgers University, Piscataway NJ 08855, USA
vincentm,parashar@caip.rutgers.edu

Abstract. The increasing scale, complexity, and interdependencies between policies and configurations of emerging networks calls for a shift from the current human-centric network management, towards more autonomic and adaptive mechanisms. Network topology represents the most fundamental network structure, and the lack of control over it has direct implications on performance, resilience, and security. This paper proposes *Topos*, an autonomic topology control approach that builds on concepts of emergence, self-organization, and graph theory to evolve and adapt the network topology to satisfy dynamically changing application requirements. Adaptation is performed in a decentralized manner where local nodes maintain neighborhood information and apply local neighbor selection policies to evolve the topology.

1 Introduction

Networked systems are continually growing in scale and complexity. For example, the last few years alone have seen exponential growth in the number of transactions, both at the network (e.g. BGP route updates and withdrawals) and at the application level (e.g. DNS, email) [1]. This rapid growth has had a significant impact on the performance, manageability and reliability of emerging networks. While the limitations induced by such a fast growth have been remedied, to a certain extent, through manual changes made to configuration tables (e.g. BGP tables), such ad hoc approaches are quickly becoming insufficient and impractical. The emerging field of Autonomic Communication aims at addressing these issues by formulating and deploying adaptive, self-configuring, and self-managing protocols and mechanisms that relieve the persistent need for human-driven management. While policies are defined by humans, their deployment, enforcement, and dependency conflict resolution should all be handled by automatic mechanisms to optimize system performance.

However, achieving this vision requires conceptual, physical and logistical modifications to existing systems and protocols.

A key issue that is significantly impacting emerging networks and applications, and can be potentially addressed by autonomic communication, is the absence of accurate knowledge of, and control over the actual topology of large networks. Network topologies define the link relationships between the nodes in the network, and have a direct impact [2] on the security, resilience, and performance of a distributed system. Large scale systems such as the Internet have grown out of minimizing cost and maximizing performance at the expense of flexibility [3–5]. For example, in the case of the Internet, it is only recently that the criticality of these concerns has been discussed, for example see *Overcoming the Internet Impasse through Virtualization* [6] and *The Internet is Broken* [7]. While it is still possible to monitor the state of the approximately 40,000 autonomous systems in existence [1] at the BGP level, monitoring all the traffic inside these autonomous systems is a daunting and nearly intractable task. It requires days of data collection and processing to obtain a global picture of the system [8, 9], which is outdated by the time it is completed. It is for these reasons, and maybe not surprisingly so, that in the past few years most distributed applications have been designed on top of, or as *overlay networks*, giving network and application engineers more control over their target network. While the physical network topology is hard-set and can not be changed using software alone, an overlay network can be rewired virtually in any desired way, thereby enabling dynamic software-driven configuration and management of the topology. Such overlay networks have been used to study problems of scalability, routing, resilience, fault tolerance, security and search in networked systems. The overall objective of the research presented in this paper is to investigate autonomic mechanisms *for gaining runtime control of the network topology on which distributed applications are deployed*. This paper proposes *Topos*, an autonomic topology control approach that evolves and adapts the network topology to satisfy dynamically changing application requirements.

The approach proposed in this paper is based on concepts of emergence, self-organization and graph theory, and has three key aspects: first, translating application requirements into desired properties of the network graph; second, introducing a metric we call *network entropy* that is compiled from the graph properties and the application requirements; and finally, adapting local neighborhoods (or subgraphs) of the entire network graph by perturbing localized subgraphs structures so as to optimize *network entropy*, and globally meet the application requirements. Adaptations are performed in a decentralized manner where local nodes maintain neighborhood information and apply local neighbor selection policies to maintain and adapt the topology. Preliminary results from proof-of-concept simulations are used to demonstrate that these localized adaptations (i.e. microconfiguration) are reflected at the level of the whole graph (macroconfiguration). Note that the absence of global knowledge distinguishes this problem from traditional graph theoretic problems that assume a globally known set of vertices and edges.

The rest of this paper is organized as follows. Section 2 refines the problem definition and presents an illustration. Section 3 presents background and related work. Section 4 describes the proposed *Topos* autonomic network topology control approach and presents some preliminary results. Section 5 presents a summary and outlines future work.

2 Problem Description

The network topology is the most fundamental network structure on which a distributed application is deployed. The lack of control over the topology has a direct impact on performance, resilience, and security of these applications. Network Engineers have to deal with networks that are physically hard or impossible to change, particularly at remote locations under different autonomous system control; what can change however, is the virtual representation of that network; two nodes physically connected by a path of length m greater than 1, can be represented in the virtual network as connected by a path of length 1. The virtualized topologies can further be used to recommend alternate physical connectivities. To illustrate the significance of network topology, consider the two canonical topologies shown in figure 1, *full* figure 1 (left) and *star* figure 1 (right). Both topologies consist of 8 vertices but have opposite topological properties. For example, in the case of *full*, communication between any two active nodes is disrupted only if all nodes fail at once, whereas in the case of *star* communication fails if the single central node fails - i.e., *full* is more resilient than *star*. On the other hand, *star* is more manageable and less expensive than *full*, which is one reason why *star* topologies are dominant in the Internet. Since applications running on emerging networks are increasingly dynamic, the most appropriate topology for these applications should also change at runtime. Consequently, autonomic and adaptive topology control is becoming a critical requirement for next-generation networks. This control involves first gathering

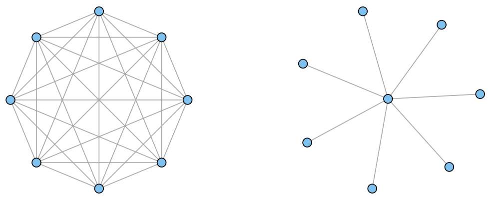

Fig. 1. Example topologies fully connected (left) and star (right)

a view of the graph structure and secondly performing connection, disconnection, or rewiring of nodes in the topology. This problem becomes exponentially hard to solve as the size of the network grows. To illustrate this point, consider structural distance, a measure of the structural differences as hamming distances between two graphs that is independent of node labeling. The structural distance between the graphs shown in figure 1 is 30, indicating that 30 operations -addition or deletion of an edge- are required to make the two graphs isomorphic. The goal of the approach presented in this paper is to develop autonomic strategies for evolving topologies to better match application requirements, based on localized (i.e. neighborhood) decisions and adaptations alone. More specifically, the problem addressed by the paper can be summarized as follows: *Given a connected network graph, what rules and adaptations applied at local neighborhood levels (i.e. microscopic) of the network lead to overall network topologies that can better support application requirements at the whole network level (i.e. macroscopic).*

3 Background and Related Work

Research efforts related to network topologies can be broadly classified as related to “topology awareness”, “topology modeling”, and “emergent topologies”, as presented in Table 1. These research areas further develop into sub-areas or applications as illustrated in figure 2. The three categories are discussed in more detail in this section. Note that the more general problem of optimizing a graph topology is inherently multi-disciplinary and has been studied in various disciplines including biological computing, artificial intelligence, graph theory, and topology modeling. While approaches based on *Parallel Distributed Processing*, *Swarm Intelligence*, and *Statistical Mechanics* offer distinct and efficient ways to address specific formulations of similar problems, the problem addressed in this paper is different from these primarily in the absence of global knowledge of the network at any time.

Application class	**Description**	**Main Result**
Topology Awareness	Use localized information to assess the current and next state of the system.	Structured or geographic-based systems can theoretically achieve logarithmic-time data propagations.
Topology Modeling	Identify best-fitting statistical models to the observed data collected from real topologies.	Large unstructured networks follow a power-law model.
Emergent Topologies	Build on the concept that the whole topology has features that are greater than the sum of its parts.	Optimal solutions to known NP-complete problems to date are provided by this approach.

Table 1. Classification of research directions related to network topology

Topology-awareness While recent research in virtualization [6] has been arguing in favor of data access protocols on the Internet that virtualize content distribution, abstracting the data from its physical location into its virtual network location doesn't necessarily serve location-dependent applications such as:

- Ad-Hoc Routing: mobile networks that can communicate only within a specific geographical space.
- Security: confidential data that should be transmitted only within trusted networks.
- Data Management: data mining and information retrieval systems that collect and catalog information based on locality.

Such applications need to maintain a proximity view of their neighbors and, as a result, are termed as *locality* or *topology-aware*. The maintenance of a correct and persistent view of a dynamic system requires

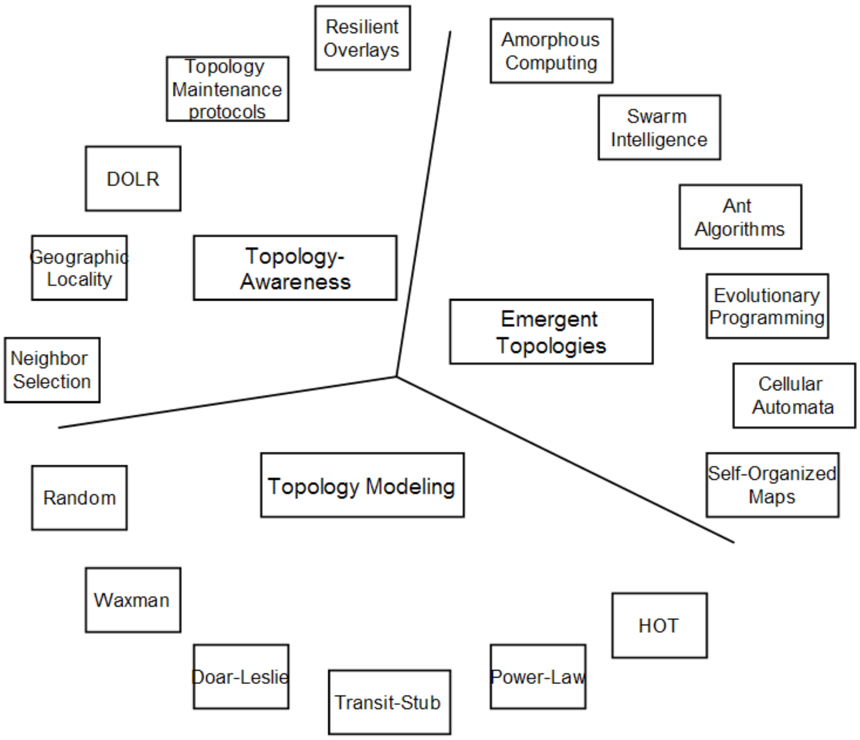

Fig. 2. Scope of research related to network topology

dedicated synchronized protocols that are either non-existent or cannot be implemented on an unreliable messaging substrate. This makes finding the appropriate node to communicate with in an ad-hoc manner a challenging problem.

Topology maintenance requires mechanisms to collect information from a set of nodes in the system in order to derive a view of the global state. The cost of maintaining a full view of the network is proportional to the size of the network [10]. While such maintenance protocols can achieve logarithmic time for infrequent changes, the number of messages exchanged grows exponentially for frequent changes.

A more recent but conceptually similar idea is to use geographic locality for topology building. Several methods have been proposed for maintaining geographic locality when a node joins a network. These include approaches based on landmark servers for position calculation [11], and on translating network distance into geographic distance [12]. However, as in the case of topology maintenance systems, these solutions do not scale well and, in the case of landmark servers, require reliable and available nodes present in the network at all time in order to extract the position of the node.

Researchers at the University of Bologna have used topology awareness to evolve a topology towards a predetermined connectivity map [13]. Their work has demonstrated the efficiency of the gossip algorithm in reaching *eventual consistency* among nodes in a large network, and have shown that their algorithm can evolve a large network of nodes towards a given topology in a few tens of cycles. While our research shares some similarities with this work in the mechanisms of topology modifications, our focus is strictly on adapting the topology to changing application requirements at runtime.

Other more recent techniques for locality and topology awareness, such as *Geographic Layout*, *Proximity Routing*, and *Proximity Neighbor Selection (PNS)* [14, 15], are evaluated in [16]. These geographic or proximity topology aware protocols determine an optimal neighbor to forward data to and build robust structured and unstructured networks. However, the evolution and adaptation of the topology is not policy-driven and can not be used to dynamically optimize the overall network.

Topology awareness has also been used for fault tolerant routing. For example, the Resilient Overlay Network [4] addresses problems due to autonomous system routing failures by proposing an overlay network protocol that dynamically determines a new route for the packets routing through a faulty node. While RON was shown to successfully reroute data around a fault, it does not attempt to reorganize the topology in order to optimize application-level functionality. Similar issues are also encountered in peer-to-peer networks, in which unreliable communications are established between peers in the absence of centralization. One such class of P2P application, *Distributed Object Location and Routing* (DOLR) [17], uses *locality awareness*, *proximity routing*, *data replication*, and *soft-state maintenance* techniques to ensure reliable and high performance search in the P2P system.

Topology-awareness has also been used in mobile ad-hoc networks (*MANET*). The cost of routing in MANETs largely influences the power consumption of the interacting devices. Each node in a MANET can maintain a local view of its neighboring nodes, i.e., a local map of the topology that, when optimized, can improve the performance of the application and the power consumption of the nodes. One such optimization is used in Ascent [18], an energy-saving protocol for sensor networks. Ascent uses local topology information and the density of packet loss to determine the node's current and next state and uses Directed Diffusion [19] to build a global view of the topology.

Network Topology Modeling In contrast to the approaches presented above that grow a network in a topology-aware manner, the network modeling field looks specifically at real-world networks, and attempts to extract a characteristic model that best fits the observed data. For example, five different models for the Internet, pure random, Waxman, Doar-Leslie, Exponential, and Locality, are presented and compared in [20]. The Internet topology was originally believed to be a random graph [21] and follow a Poisson distribution. In 1999 several independently conducted studies [5, 22, 3] demonstrated that the distribution of router links on the Internet exhibited power-law distribution. The shared content on the Internet was shown in [22] to be a *scale-free network*. More recent studies have elaborated on these initial observations, always confirming the power-law distribution but giving different explanations for the manifestation of these properties [23]. A study of the evolution of the Internet topology design is presented in [23], where the study of a single Internet Service Provider is used to extrapolate information of the global Internet. Network topology modeling is fundamental to understand the advantages and drawbacks of existing networks, as well as defining appropriate metrics. This research applies results from topology modeling towards topology optimization and self-configuration to improve application performance.

3.1 Emergent and Bio-Inspired Approaches

Emergence, or the emergent property, is a characteristic of complex systems, where the number of possible interactions is so large that the system appears greater than the sum of its parts. For example, the World Wide Web exhibits emergent properties, as links under no centralized control follow a power-law distribution. There are several research efforts that build on the concept of emergence to construct and manage network topologies [24, 13]. While the research presented in this paper is also based on emergence, it differs in the strategies used to adapt node connectivity and evolve the topology, as well as the metrics used to evaluate candidate adaptations.

Other related efforts that are inspired by nature and biological systems [25] include Amorphous Computing [26], Swarm Intelligence [27, 28], and Cellular Automata [29]. Amorphous Computing applies concepts from biology and evolution to develop computer languages (Growing Point Language [30]) and decentralized evolving systems [31]. Each computing cell is viewed as equivalent to an organic cell that is guided by its environment in order to determine its next state. The cells follow virtual chemical gradient and density trails to identify their position and direction of evolution. Swarm Intelligence is inspired by social and behavioral theories in the animal and human kingdom. For example, foraging, nest building, and burial activities of social insects (ants, wasps, termites) follow local rules applied by each entity with no knowledge of the whole. Similarly flocks of birds, schools of fish, herds of mammals, follow local rules that lead to structures with emergent properties. Cellular automata is a mathematical model that was introduced by Jon Von Neumann and originated with the theory of computing. It is characterized by the evolution of a set of self-replicating nodes with deterministic rules that result in a random global behavior.

The strategies discussed above are similar in that they are based on the manifestation of certain global properties from the interaction of small and simple parts. As a result, these approaches are inherently nondeterministic and their quantification a key challenge. Existing research efforts in this direction raise important questions about the evolution of complexity of networks, what distributions are manifested, how to quantify their resilience, security, and efficiency, and what rules of evolution to use to obtain a desired behavior? Answering these questions is what motivates our work.

4 Topos: An Approach to Autonomic Topology Control

This research applies concepts of emergence and self-organization to the autonomic control of network topologies. A key challenge in approaching this problem is defining the metrics and deriving the rules to be used by the nodes to determine the nature and direction of evolution. To address this problem, we propose the concept of *network entropy*, which is a compound metric based on the measure of the number of feasible configurations that a topology can take for a given application context.

The overall three-step approach is illustrated in figure 3 and is described in the rest of this section.

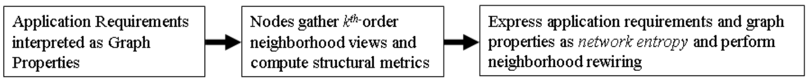

Fig. 3. High-level view of *Topos* Approach

Application Requirements: An essential characteristic of distributed applications is that they have to deal with the communication uncertainty and unreliability as part of their design requirement. When not handled appropriately, errors in the physical, network, or transport layer lead to failure of the application to meet its runtime objective. While an application requirement can involve a very large number of parameters, such as performance, cost, message format, level of service, the network aspects of these application requirements can be expressed using Quality of Service traffic engineering measures such as delay, jitter, loss, order, and error. Each of these parameters can be correlated with properties of the underlying network graph, and correspondingly, the application's QoS requirements and constraints can be translated into requirements and constraints on the topology of the network.

Graph Metrics and their correlation with application requirements: Table 4 lists the graph metrics that are used in *Topos* to gather neighborhood information at each node. These metrics constitute the building blocks from which *network entropy* and the application-dependent evolution rules are derived. Most of these metrics are presented and discussed in length in [32, 33]. The metric can be classified as local or global, depending on the number of nodes involved in computing the metric. A metric is local when a single node derives the metric, and global when the whole graph is considered to obtain the metric.

The metrics can be further categorized based on the type of result that they produce. For a graph with N vertices, the metrics *degree, clustering coefficient, betweenness, spectrum, and eigenvector centrality*, produce a list of results of dimension N; the metrics *degree distribution, joint-*

Local Metrics	**Summary**	**Range**
Degree	number of neighbors	$[0, n-1]$
Average Degree	average number of neighbors per node over all nodes	$[0, n-1]$
Average Clustering	how close to fully connected is the graph	$[0, 1]$
Global Metric	**Summary**	**Range**
Degree Distribution	probability distribution of node degrees	$[0, 1]$
Joint Degree Distribution	distribution of average neighbor degree of a node of average degree d	$[0, maxdegree]$
Diameter	longest path between any two nodes	$[1, n-1]$
Betweenness	number of times a node is present in the shortest path between every pair of nodes	$[0, 1]*$
Smetric	normalized sum of product of degrees of nodes connected by a link	$[0, 1]*$
Spectrum	set of eigenvalues of the adjacency matrix of the graph	$[-n, n]$
Eigenvector Centrality	scoring eigenvectors by importance	$[-1, 1]$
Assortativity	measures the preference of a node to connect to like degrees	$[-1, 1]*$

Table 2. Network Metrics. * indicates normalized values.

degree distribution produce a list of dimension in the range $[1, n-1]$, limited by the maximum node degree in the network; and finally the metrics *smetric, assortativity, average degree, and diameter* produce a single number that captures global properties of the graph.

Figure 4 illustrates a few of these metrics for canonical topologies. For example, consider the two topologies presented in Section 2, i.e., fully connected and star networks. Their associated metrics are presented in Table 3. These metrics support the observations made in Section 2 regarding relative resilience properties of *star* and *full*,. The *betweenness* metric clearly shows that *star* has a single central node through which all traffic flows, whereas *full* has all nodes at an equal betweenness value. Further, the *joint degree distribution* metric shows that for *star* the average neighborhood degrees of nodes of average degree 2 is highest, while for *full* it is unique and highest for nodes of degree 7. This indicates the centralized structure of *star* versus the resilience of *full*. The *spectrum* metric is yet another indication that *full* is better connected than *star* with a larger second eigenvalue, also known as the em fiedler value.

Network Entropy: As mentioned above, *Topos* introduces the concept of *network entropy* that can be effectively used to guide and con-

Metric Name	full	star
Number of Nodes	8	8
Number of Edges	28	7
Degree	7 7 7 7 7 7 7 7	7 1 1 1 1 1 1 1
Clustering	1	0
Betweenness	0 0 0 0 0 0 0 0	21 0 0 0 0 0 0 0
Spectrum (2 largest)	7 -1	2.645751 0
Eigenvector Centrality	0.3535534 0.3535534	0.7071068 0.2672612
Degree Distribution	0 0 0 0 0 0 0 1	0 0.875 0 0 0 0 0 0.125
Joint Degree Distribution	0 0 0 0 0 0 1	1 0 0 0 0 0 1
Average Degree	7	1.75
Diameter	1	2
Smetric	1	0.28
Mixing Pattern	NA	-1

Table 3. Summary of graph metrics for the fully connected and start topologies.

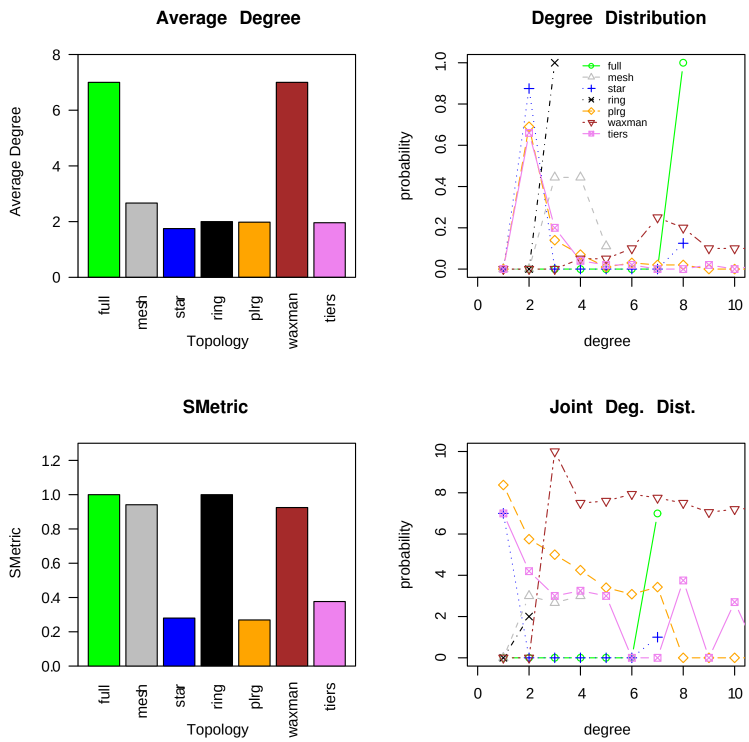

Fig. 4. Graph metrics for fully connected and star topologies

strain local adaptation at individual network nodes. From a structural approach the graph connectivity set is a permutation over all possible edges between all pairs of nodes. For a graph of size N, this ranges from a set of edges of size $O(N)$ to a fully connected network of size $O(N^2)$. The most appropriate configuration depends on application requirements and on the current state of the network nodes. *Network entropy* is a compound metric that is derived from the graph metrics and the application constraints on these metrics. It represents the correlation between the application requirement and the local graph properties, and is used to evaluate microlevel connectivities and select adaptation strategies.

Morphing micro to macro configurations: *Topos* considers the global network as a collection of local neighborhoods and adapts local neighborhoods so as to minimize network entropy, which in turn leads to global network with respect to an application metric.

4.1 A Proof-of-concept Example

Consider the network presented in figure 5 (a). A small fraction of the total number of nodes are selected to build a second-order neighborhood view. This view includes each node's neighbors as well as its neighbors' neighbors. These three neighborhoods are presented in figure 5 (a-67, a-23 and a-12). We will refer to the whole graph as the *macrograph* and the neighborhood subraphs as *micrographs*. The *macrograph* is a scale-free network that exhibits a power-law degree distribution. A summary of the metrics for this network is presented in Table 4 (Initial Topology column).

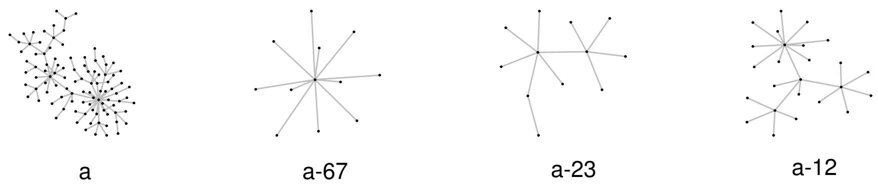

Fig. 5. Initial topology: Scale-Free Network (leftmost) and 3 of its neighborhoods

Consider an application that requires a resilient network with a low number of intermediary hops. These requirements set constraints on the

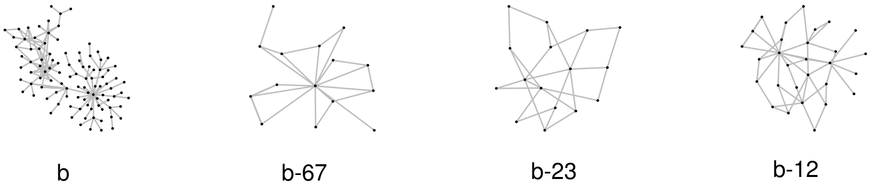

Fig. 6. Perturbed topology: Scale-Free Network (leftmost) and 3 of its neighborhoods.

Metric Name	**Initial Topology**	**Perturbed Topology**
Average Betweenness	177	145
Diameter	11	9

Table 4. Summary of metrics for the proof-of-concept simulation.

betweenness and diameter of the entire graph. A direction for adaptation then is to modify the subgraphs such that the values for betweenness and diameter decrease. Using network entropy to measure the state of the graph with respect to these desired properties, we rewire the nodes in each subgraph such that lower degree nodes are more connected and some highly connected nodes removed. The betweenness and diameter of this perturbed topology is presented in Table 4 (Perturbed Topology column) and shows that average betweenness and diameter of the overall topology improved by 18%. Using this proof-of-concept simulation, only 6% of the subgraphs and second-order neighborhood views, the overall network is more resilient (a higher clustering coefficient), has an improved mixing pattern (characteristic of robust networks), and presents a lower overall betweenness value and a lower overall diameter. Note that the topology graphs were produced and the simulations conducted using the *R* [34] statistical software package and the *igraph* contributed package.

5 Summary and Future Work

The increasing scale and complexity of networks has necessitated autonomic communication approaches that provide engineers and application developers with more control of the network on which their applications are deployed. Since network topology has a direct impact on performance, resilience, and security, autonomic control of the network topology is an important and immediate requirement. This paper proposed *Topos*, an autonomic topology control approach that builds on concepts of emergence, self-organization and graph theory. *Topos* evolves and adapts the network topology in an autonomic manner to satisfy dynamically changing application requirements. Adaptation is performed in a decentralized manner where local nodes maintain neighborhood information and apply local neighbor selection policies to evolve the topology. Using prototype simulations of *Topos*, this paper demonstrates that by modifying the network at the microscopic level (neighborhood level) the macroscopic level (global topology) is improved in terms of applications-centric performance metrics.

Ongoing and future work focuses on several outstanding aspects of *Topos* formulation including: establishing requirements for deploying and enforcing topology rewiring policies; addressing conflict detection and resolution challenges that arise when rewiring involves nodes that are a part of different neighborhoods; the analysis and characterization of various evolutionary strategies, and the evaluation of their applicability and benefits for different categories of application requirements.

References

1. G. Huston, "The bgp report," 2006.
2. K. P. Gummadi, R. Gummadi, S. D. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica, "The impact of dht routing geometry on resilience and proximity," in *Proceedings of the ACM SIGCOMM 2003*, 2003.
3. J. M. Carlson and J. Doyle, "Highly optimized tolerance: A mechanism for power laws in designed systems," *Physics Review E*, vol. 60, no. 2, pp. 1412–1427, 1999.
4. D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proceedings of the eighteenth ACM symposium on Operating systems principles*, pp. 131–145, ACM Press, 2001.
5. M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *SIGCOMM*, pp. 251–262, 1999.
6. T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the internet impasse through virtualization," *Computer*, vol. 38, no. 4, pp. 34–41, 2005.
7. D. Talbot, "The internet is broken," *Technology Review*, January 2006 2006.
8. CAIDA, "Cooperative association for internet data analysis," 2006.
9. ISC, "Internet systems consortium," 2006.
10. H. Abu-Amara, B. A. Coan, S. Dolev, A. Kanevsky, and J. L. Welch, "Self-stabilizing topology maintenance protocols for high-speed networks," *IEEE/ACM Trans. Netw.*, vol. 4, no. 6, pp. 902–912, 1996.
11. T. S. E. Ng and H. Zhang, "Towards global network positioning," in *Proceedings of the First ACM SIGCOMM Workshop on Internet Measurement*, pp. 25–29, ACM Press, 2001.
12. M. Costa, M. Castro, A. Rowstron, and P. Key, "Pic: Practical internet coordinates for distance estimation," in *Proceedings of the International Conference on Distributed Computing Systems (ICDCS'2004)*, 2004.

13. M. Jelasity and O. Babaoglu, "T-man: Gossip-based overlay topology management," in *Third International Workshop on Engineering Self-Organising Applications (ESOA'05)*, (Utrecht, The Netherlands), 2005.
14. D. R. Karger and M. Ruhl, "Finding nearest neighbors in growth-restricted metrics," in *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pp. 741–750, ACM Press, 2002.
15. M. Castro, P. Druschel, Y. Hu, and A. Rowstron, "Topology-aware routing in structured peer-to-peer overlay networks," Tech. Rep. MSR-TR-2002-82, Microsoft Research, One Microsoft Way, Redmond, WA 98052, 2002 2002.
16. M. Castro, P. Druschel, Y. C. Hu, and A. Rowstron, "Exploiting network proximity in distributed hash tables," in *Proceedings of the International Workshop on Future Directions in Distributed Computing (FuDiCo 2002)*, 2002.
17. B. Zhao, A. Joseph, and J. Kubiatowicz, "Locality aware mechanisms for large-scale networks," in *In Proceedings of Workshop on Future Directions in Distributed Computing*, 2002.
18. A. Cerpa and D. Estrin, "Ascent: Adaptive self-configuring sensor networks topologies," *IEEE Transactions on Mobile Computing*, vol. 03, pp. 272–285, 2004.
19. C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000)*, 2000.
20. E. Zegura, K. Calvert, and S. Bhattacharjee, "How to model an internetwork," in *Fifteenth Annual Joint Conference of the IEEE Computer Societies, Networking the Next Generation.INFOCOM '96.*, vol. 2, pp. 594–602, 1996.
21. P. Erdos and A. Renyi, "On random graphs," *Publ. Math. Debrecen 6*, vol. 6, pp. 290–297, 1959.
22. A.-L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, 1999.
23. D. Alderson, J. Doyle, R. Govindan, and W. Willinger, "Toward an optimization-driven framework for designing and generating realistic internet topologies," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 1, pp. 41–46, 2003.
24. B. F. Cooper and H. Garcia-Molina, "Ad hoc, self-supervising peer-to-peer search networks," *ACM Trans. Inf. Syst.*, vol. 23, no. 2, pp. 169–200, 2005.
25. N. Forbes, *Imitation of Life: How Biology is Inspiring Computing*. MIT Press, 2004.
26. H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T. F. Knight, Jr., R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss, "Amorphous computing," *Commun. ACM*, vol. 43, no. 5, pp. 74–82, 2000.
27. R. C. Eberhart, Y. Shi, and J. Kennedy, *Swarm Intelligence*. Morgan Kauffmann, 2001.
28. E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence From Natural to Artifical Systems*. Oxford Press, 1999.
29. S. Wolfram, *A new kind of science*. Champaign, Ilinois, US, United States: Wolfram Media Inc., 2002.
30. D. N. Coore, *Botanical computing: a developmental approach to generating interconnect topologies on an amorphous computer*. PhD thesis, MIT, 1999.
31. R. Nagpal, "Programmable self-assembly using biologically-inspired multiagent control," in *Proceedings of Autonomous Agents and Multiagent Systems (AAMS 2002)*, 2002.
32. S. Dorogovtsev and J. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, 2004.
33. P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. Claffy, and A. Vahdat, "Lessons from three views of the internet topology," Tech. Rep. CAIDA-TR-2005-02, CAIDA, 2005.
34. "R-project." http://www.r-project.org, 2006.