Towards A Theory of Autonomous Reconstitution of Compromised Cyber-Systems

Pradeep Ramuhalli[1], Mahantesh Halappanavar[1], Jamie Coble[2], Mukul Dixit[1]
[1]Pacific Northwest National Laboratory, Richland, WA USA
Pradeep.Ramuhalli@pnnl.gov, hala@pnnl.gov, Mukul.Dixit@pnnl.gov
[2]The University of Tennessee, Knoxville. TN USA
jcoble1@utk.edu

***Abstract*—Effective reconstitution approaches for cyber systems are needed to keep critical infrastructure operational in the face of an intelligent adversary. The reconstitution response, including recovery and adaptation, may require significant reconfiguration of the system at all levels to render the cyber-system resilient to ongoing and future attacks or faults while maintaining continuity of operations. A theoretical basis for optimal dynamic reconstitution is needed to address the challenge of ensuring that dynamic reconstitution is optimal with respect to resilience metrics, and is being developed and evaluated in this project. Such a framework provides the technical basis for evaluating cyber-defense and reconstitution approaches. This paper describes a preliminary framework that may be used to develop and evaluate concepts for effective autonomous reconstitution of compromised cyber systems.**

Keywords-resilient cyber-systems; reconstitution; recovery; evolution

I. INTRODUCTION

The ability to maintain mission-critical operations in cyber-systems in the face of disruptions is critical. Faults in cyber systems can come from accidental sources (e.g., natural failure of a component) or deliberate sources (e.g., an intelligent adversary). Natural and intentional manipulation of data, computing, or coordination are the most impactful ways that an attacker can prevent an infrastructure from realizing its mission goals. Under these conditions, the ability to reconstitute critical infrastructure becomes important. Specifically, the question is: *Given an intelligent adversary, how can cyber systems respond to keep critical infrastructure operational?*

In cyber systems, the distributed nature of the system poses serious difficulties in maintaining operations, in part because a centralized command and control apparatus is unlikely to provide a robust framework for resilience. Resilience in cyber-systems, in general, has several components, and requires the ability to anticipate and withstand attacks or faults, as well as recover from faults and evolve the system to improve future resilience. The recovery effort and any subsequent changes may require significant reconfiguration of the system at all levels—hardware, software, services, permissions, etc.—if the system is to be made resilient to further attack or faults. This is especially important in the case of ongoing attacks, where reconfiguration decisions must be taken with care to avoid further compromising the system while maintaining continuity of operations. Collectively, we will label this recovery and evolution process as “reconstitution.” Currently, reconstitution is performed manually, generally after-the-fact, and usually consists of either standing up redundant systems, check-points (rolling back the configuration to a “clean” state), or re-creating the system using “gold-standard” copies. For enterprise systems, such reconstitution may be performed either directly on hardware, or using virtual machines.

A significant challenge in this context is the ability to verify that the reconstitution is performed in a manner that renders the cyber-system resilient to ongoing and future attacks or faults. Fundamentally, the need is to determine optimal states of the cyber system when a fault is determined to be present.

Contributions: This paper presents preliminary research towards concepts for effective autonomous reconstitution of compromised cyber systems. We describe a *mathematical formulation* as a first step towards a theoretical basis for developing and evaluating autonomous reconstitution algorithms in dynamic cyber-system environments. We then propose formulating autonomous reconstitution as an *optimization problem* and describe some of the challenges associated with this formulation. This is followed by a brief discussion on potential solutions to these challenges.

II. BACKGROUND

The quest for resilient cybersystems has seen a number of potential approaches, each of which attempt to add specific properties to the system that make it resilient to one or more attack vectors. Examples of these properties include diversity, redundancy, deception, segmentation, and unpredictability. One approach that incorporates some of these elements is moving target defense [1]. These types of methods are usually heuristic and can be shown using empirical approaches to provide some level of resilience. However, these approaches cannot be readily applied after a system has been compromised. It is also not clear whether such methods are

978-1-4799-1535-4/13/$31.00 ©2013 IEEE

applicable under any attack scenario, or if there are specific limitations. Fundamental to understanding the applicability of many of these approaches as attacks evolve is the ability to mathematically define cyber-systems in some manner. Specifically, there is a need to determine the key properties of a cyber-system, determine nominally safe configurations of the system in terms of one or more metrics, and define the mathematical framework that uses this information.

As might be expected, the problem of cyber-system resilience has seen significant research over the past few years. A number of groups (MITRE, Raytheon) have proposed frameworks for resilience that encompass the basic ideas (anticipate, withstand, recover, evolve) (Fig. 1). The right side of the figure (anticipate, withstand) might be generally thought of as enabling robust design of the cyber-system, while left-half corresponds to reconstitution in the event of compromise. However, these two elements (robustness and reconstitution) are not independent, and leverage information from each other (for instance, information available during reconstitution may inform robust design, and reconstitution approaches may be constrained by robust design concepts).

Frameworks for resilience that incorporate operational aspects have also been proposed. Reference [2] also proposes an operational framework for resilience but acknowledges the difficulties in a practical implementation. Reference [3] discusses the use of probabilistic risk assessment as a tool for understanding and improving resilience.

III. Related Work

Prior work on resilience (and recovery from attacks) can largely be categorized into approaches based on fault-tolerance algorithms in a Byzantine fault environment [4], and approaches based on moving target defense [1, 5]. While existing theories for fault tolerance (e.g., Byzantine fault tolerance) can guarantee resilience under certain conditions [6, 7], in practice, these theories can break down in the face of an intelligent adversary. Further, it is difficult, in a *dynamically* evolving environment, to determine whether the necessary conditions for resilience have been met, resulting in difficulties in achieving provably resilient operation. In addition, existing theories often do not sufficiently take into account computational cost [8, 9] (adversary is assumed to have infinite resources and time), hierarchy of importance (all network resources are assumed to be equally important), and the dynamic nature of some attacks (i.e., as the attack evolves, can fault tolerance be maintained?).

A number of other research developments may be of relevance. These include self-stabilizing systems [10, 11], distributed algorithms in systems with sectional faults [12], and self-organizing systems [13]. Each of these approaches has the potential to improve cyber-resilience. However, these theories will need to be augmented to account for an intelligent adversary. Conversely, game theory and other conflict models [14] bear on intelligent adversaries, but may not always account for faults.

Recent publications (such as [15]) indicate that a dynamic reconstitution and reconfiguration effort may be capable of addressing certain classes of attacks. However, the published information indicates a manual, "operator-in-the-loop" approach, where, once indicators of compromise are identified, operators at a Resilience Operations Control System and the Cyber Operations Center decide on an appropriate course of action and implement it. However, the test implementation of their framework for resiliency appears to have depended on commercial resiliency management systems, the market for which was fairly immature as of the writing of that article. Products in R&D phase were stand-alone and were not easily integrable with enterprise-wide resiliency management systems, and commercially available systems were generally based on minor modifications to existing security products that did not meet the needs for resiliency. Further, the work does not appear to focus on asymmetry. Some prototypical products that have been discussed in the literature include the Net Maneuver Commander [16].

This paper describes a *state-space* based formulation for use in reconstitution of cyber-systems. Specifically, the state C_t of a cyber-system at time t is a representation of its key properties. When a system is compromised, it moves from a fully operational state to one of several compromised states (Fig. 2) where it loses the capacity to maintain continuity of operations. The reconstitution effort is then one of moving the system back to one of several fully operational states possibly through several intermediate states. This process may be defined as one of optimization, with metrics for resilience and continuity of operations used to determine, at each time step subsequent to the compromise, the optimal state (such as network connectivity, services, and hosts) to ensure continuity of operations while improving resilience.

IV. Mathematical Preliminaries

We use C_t to represent the state of a cyber-system at time t. In this paper, we assume that a graph G_t may be used to represent the network connectivity information within the system at time t. A graph $G = (V, E)$ is a pair, where the set of vertices V represents unique entities in a system, and the set of of edges E represents binary relations between vertices. The connectivity information may change over time as the cyber-system goes through the reconstitution process, and this dependency is captured explicitly. Further, we assume that a tensor ϕ_t can be used to represent the configuration at time t of the different elements in the cyber-system. We define

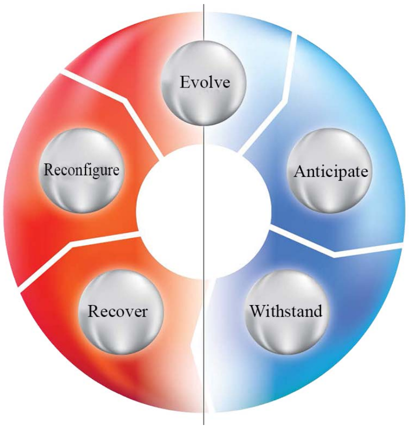

Figure 1. Elements necessary for resilience in cyber-systems

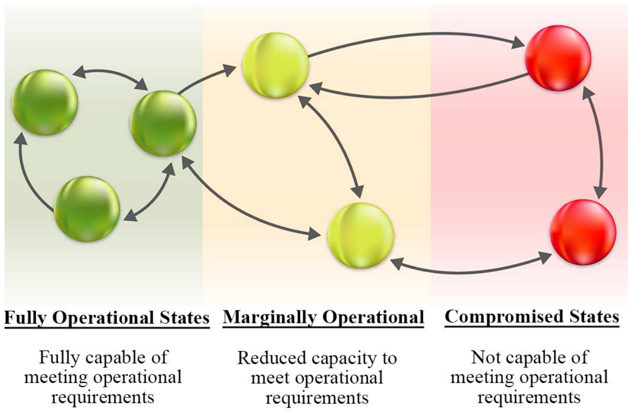
Fully Operational States
Fully capable of meeting operational requirements
Marginally Operational
Reduced capacity to meet operational requirements
Compromised States
Not capable of meeting operational requirements

Figure 2. Conceptual representation of reconstitution in cyber-systems as moving the system from one of several compromised states to a fully operational state

$$C_t \triangleq \{G_t, \phi_t\} \tag{1}$$

Note that, in general, a cyber-system is defined in continuous time; that is, connectivity and configuration information is present at all times. However, the state of the cyber-system will be assumed to be discrete; that is, it can take on only one of a finite number of values. With this constraint, the system C_t is an example of a continuous time, discrete-valued system [17].

Normal or abnormal traffic within the cyber-system is assumed to generate data D_t that is a function of C_t, and some parameters Θ:

$$D_t = f\left(C_t, \Theta\right) \tag{2}$$

The parameters Θ might represent, for instance, the portions of the cyber-system where sensors are deployed. Alternatively, Θ might be used to represent a set of transformation parameters employed to extract relevant data from the cyber-system. In any case, (2) is generic enough to cover many possibilities. For the moment, we place no restrictions on whether this data is available for access from outside the system, or only within the system. We however assume that D_t is a random process defined by a probability density function on Θ:

$$\Theta \sim P(\theta) \tag{3}$$

where θ represents the parameters of the density function. This assumption is simply a reflection of the fact that, in general, the only kinds of system-wide information that may be obtainable are statistical quantities such as, for instance, the mean and variance of network traffic flow patterns, and that minute-to-minute specification of data within a cyber-system is, in general, difficult to obtain. Note that this assumption is not restrictive—in cases where a deterministic description may be available, it may be accounted for by setting the corresponding probability to 1 (interpreted as complete knowledge). We may further assume that D_t is ergodic in the mean (i.e., time averages may be used to approximate process means). This is a simplifying assumption, and will need to be validated using data from realistic cyber-systems.

Faults within the cyber-system are described within this formulation through a mathematical description of their effects. This allows for a framework that can capture both natural and adversarial faults, and enables the resulting reconstitution approach to be agnostic to specific attack vectors. We denote the fault sequence by the vector

$$F = \left[F_1, F_2, ..., F_t, ...\right] \tag{4}$$

where F_t denotes the fault at time t, and F_k, $1 \le k \le t-1$ denote the sequence of faults before time t; note that future faults (i.e., for times greater than t) may be incorporated into this framework. As noted earlier, these faults may be either natural or due to an intelligent adversary.

Faults at any time instant may result in hardware level effects (such as loss of a server) and result in a change in the connectivity, and consequently, in changes in the configuration and data. Alternatively, faults may manifest themselves at only the configuration level (for instance, a change in the firewall settings on a specific server), or within the data available in the cyber-system (for instance, higher than normal numbers of open ports, or increased network traffic). Thus, we may assume that, at time t, fault F_t defines a mapping of the form:

$$F_t : C_t \to C_{t+1} \tag{5}$$

The resulting state is generally assumed to be a less-desirable state (from the perspective of maintaining mission-critical operations, and as defined by some metrics—see below). For simplicity, if no fault (natural or otherwise) exists in the system at a given time (say $t = k$), we denote $F_k = 0$.

We use the representation above to define the needs for reconstitution. We do not claim that this representation is unique, and as other representations of cyber-systems, data, and fault sequences become available, may be easily incorporated in this work.

V. Metrics for Continuity of Operations

Critical to reconstitution is the ability to define continuity of operation. In this initial formulation, we assume that a metric exists that can be used for this purpose. Such a metric M_t would need to be a function of the system state:

$$M_t = g\left(C_t\right) \tag{6}$$

and must be computable from knowledge of the system network topology and configuration information. A number of resilience metrics related to continuity of operations are defined in the literature [18]. However, the bulk of these metrics are focused on system-level quantities (such as time to recover from an attack, percentage of available services, etc.). While these are important and help characterize the system performance, these are difficult to use for dynamic reconstitution, as computing such metrics in real-time (as the system is being reconstituted) from knowledge of only the configuration and/or connectivity is difficult. Instead, indirect metrics are necessary, and include graph metrics such as diameter, algebraic connectivity, average path length, clustering coefficient, although other graph statistics may be relevant and computable in real-time. In addition to graph metrics, metrics such as vulnerability scores [19] that may be computed from configuration information as well as a priori knowledge about attack vulnerabilities may also be applicable.

An important element that needs to be captured in the metric is the potential dependencies between critical services, and between the deployment of services and the underlying network topology. This results from the need to ensure that the impact of any fault or attack and the subsequent reconstitution is represented correctly. For instance, service A and B may depend on a higher-level service C (for example, an authentication service). The impact of a fault that results in C being unavailable will impact the availability of A and B, and therefore the ability to meet mission needs.

VI. A Proposed Framework for Reconstitution

We define reconstitution as determining the state C_t, $t > T_0$ that ensures continuity of operation. Here, T_0 is assumed to be the time at which the reconstitution effort is initiated. In general, the problem of determining secure configurations is NP-complete [19]. However, solutions that are "good-enough" may still be possible in a reasonable time-frame.

One approach to formulating the concept of asymmetric resilience is by accounting for the cost to the defender [20]. Assume that the cost to the attacker can be represented by R_a while the cost to the defender is represented by R_d. The cost R_d is a dimensionless number that accounts for the infrastructure cost R_{di} during the reconstitution effort (normalized to the initial cost incurred during the original system setup) as well as the risk (of continued attack or faults) associated with the configuration. Note that R_{di} may be estimated from the connectivity graph as well as any redundancies required for system robustness. R_a may consist of similar quantities; however, measuring R_a is a difficult proposition. It is, however, reasonable to assume that the faster the reconstitution effort, the greater the cost to the attacker (both in terms of infrastructure cost to maintain an attack as well in terms of risk of attribution). For this initial framework, we will assume that R_a is inversely proportional to time spent in the reconstitution effort:

$$R_a \propto 1/(t - T_0), \quad t > T_0. \tag{7}$$

A. Reconstitution as an Optimization Problem

Given this information, one possible approach to reconstitution is to frame the problem as a multi-objective optimization problem, where asymmetric advantage may be introduced by, for instance, requiring that the solution minimizes the cost to the defender while maximizing the estimated cost to the attacker. Specifically, we want to maximize M_t and R_a while minimizing R_d by adjusting the network connectivity as well as the configuration information. This may be represented as

$$\max_{G_t, \phi_t} M_t, \frac{R_a}{R_d} \tag{8}$$

subject to constraints on graph connectivity, allowed configurations, and available resources. This framework allows for the incorporation of costs into the reconstitution effort, as well as attacks or faults during the reconstitution effort.

Potential approaches to optimization that are applicable within this framework include genetic algorithms [19], linear programming, and dynamic programming [21]. The framework also encompasses both natural and adversarial faults.

Challenges in this context include optimization and control of large-scale systems, and in the face of continuing attacks. An additional challenge is presented by the distributed nature of cyber-systems, in part because a centralized command and control apparatus is unlikely to provide a robust framework for resilience, and the reconstitution process will need to be managed in a distributed manner.

A further challenge is the ability to explicitly reduce cost to the defense. The cost to the defender, while notionally simple to calculate, is difficult to quantify during the course of optimization. This is because any change in the system state results in incurred costs to the defense; thus, the overall cost over the entire reconstitution process will be the sum of the individual costs at each step in the process. However, conventional optimization approaches focus on minimizing only the incremental cost over the next iteration in the optimization. Thus, tools for optimization that find the overall least-cost trajectory in state space are needed. Moving to a distributed approach to reconstitution adds further challenges in this regard.

In general, the assumption in this formulation is that the cost functions and constraints define a non-convex problem, and the resulting solution will only be locally optimal (and dependent on the starting state). This is because, in the context of the problem being studied here, a globally optimal solution refers to a resilient state that is capable of providing continuity of operations regardless of the fault or attack vector. While this may be theoretically possible, practical bounds on resources and time to recovery will necessarily restrict the space of possible solutions that can be explored, and result in tradeoffs.

The formulation specified in equation (8) is, in general, an example of a multi-objective optimization problem. In general, given the specific metrics cited in this paper, it is expected that a single optimal solution will be possible. Instead, the likelihood is that the solution will be a family of pareto-optimal solutions that offer varying tradeoffs between the different metrics.

A specific example of the formulation is for optimal allocation of services assuming that the underlying topology is constant. In this context, we have a set of n critical services $S_1, S_2, \ldots S_\mathrm{n}$ and assume that a reconstitution effort is complete only when all the n services become operational. The resulting configuration may be represented by the optimal ϕ^*. A number of service providers (e.g., computer servers) are assumed available, and connected in some topology. Resource constraints are assumed at each service provider, and restrict the number of possible services (defined by the subset $S_p = \{S_{p1}, S_{p2}, \ldots S_{pk}\}$) that can be hosted on server p. We are also given a cost function R_{dp} to make a service provider p operational: $R_{dp} : S_p \to \mathrm{R}^+$. We define

$$R_d = \sum_p R_{dp} . \tag{9}$$

Finally, dependencies between services are incorporated through constraints on the network topology (to minimize delays) and the configuration ϕ.

Given this setup, the goal of reconstitution is then to determine a minimum cost set of providers (each running a subset of services) such that all critical services are operational. This is a re-statement of the set-cover problem [22]: Given a universe U of n elements $S_1,S_2,\ldots S_n$, a collection of subsets of U, $S_p = \{S_{p1}, S_{p2}, \ldots S_{pk}\}$, $p = 1,2,\ldots P$ and a cost function c: $S \rightarrow \mathbf{R}^+$, find a minimum cost subcollection of S that covers all elements in U. This particular instantiation of the reconstitution problem may be addressed through applicable combinatorial optimization approaches, though the same challenges that apply to all optimization-based approaches to reconstitution apply here.

VII. Results

The proposed optimization-based approach to reconstitution was tested using a set of small cyber systems, represented as graphs, and a genetic algorithm for optimization. An n-node network with m-services is described by a graph G that captures the connectivity structure of the network. Here, a node represents some computational resource, such as a computer or a server that can run any subset of the m services. A configuration matrix, ϕ, describes the configuration of each node with respect to the critical services. The configuration of node i with respect to service j is comprised of a triplet of binary values: is service j loaded on node i, does the configuration of node i support service j, and is service j currently running on node i?

To test reconstitution, a set of random networks are simulated with random connectivity patterns and configuration matrices. Several constraints are imposed on the network to simplify the evaluation process:

- Each node can be connected to at most 10% of the total number of nodes in the network;
- Each node can only run services that are loaded on the node and supported by the node configuration;
- Each node can run at most three services; and
- Each service can be run on at most two nodes.

While many of these constraints may not be applicable in realistic cyber-systems, they nonetheless provide a starting point for evaluating the proposed framework. Evaluations with these constraints removed will be performed in the future.

Hardware faults can occur, wherein a node is permanently removed from the network and its services are lost. The genetic algorithm attempts to reconstitute the network to restore critical services and improve network resilience, subject to the same constraints defined above. The genetic algorithm's fitness function attempts to maximize network robustness, characterized by the algebraic connectivity [23], and the availability of critical services, while minimizing the cost to update the network. Costs are incurred by adding or removing network connections, loading services on nodes, and changing node configurations to support new services.

In this initial implementation, the following assumptions are made:

- Dependencies between services, or services and the underlying topology, are ignored.
- Incrementally minimizing the cost at each iteration is assumed to minimize the overall cost.
- All network nodes are assumed to be functionally equivalent.
- Any node that fails is assumed to be permanently removed from the network. As a result, the risk of reconstituted node(s) being compromised again is not considered.

These assumptions will be relaxed as this formulation is developed further.

The efficacy of this approach is demonstrated on a set of randomly generated 30-node and 40-node networks with 50 critical services running on random set of nodes. On the 30-node network, two cases are considered: first, a single fault occurs at time zero and the network is reconstituted; second, multiple faults occur, one at time zero and others as the network reconstitution is underway. A fault may result in the loss of one or more nodes. Due to space restrictions, we present results on the 30-node network with two faults (at time step zero and time step five), and the 40-node network with a single fault at time step zero. For each of these cases, we first present a rendering of the network structure (connectivity) and show the process of reconstitution with respect to algebraic connectivity and the number of services active for a given time step. Each experiment is run for 10 instances where the connectivity pattern is retained, but with different nodes randomly chosen as faults. The error bars show the variance between different instances for a given case.

Fig. 3 is a rendering of the 30-node network at time step zero with a fault resulting in the loss of three nodes. Fig. 4 and Fig. 5 show the reconstitution process in progress. While Fig. 4 shows the algebraic connectivity on the Y-axis, Fig. 5 shows the number of services. We only provide a plot of close-up for the first few time steps of recovery. It can be observed that the network services recover quickly after a fault at time step zero. But the recovery after a second fault at step five is relatively slower. Similar results for the 40 node network are presented in Figures 6-8. These results indicate that the proposed framework for reconstitution is a viable one. While not evaluated at this stage, the framework is expected to provide a means to develop new reconstitution algorithms as well as evaluate options for reconstitution through the choice of appropriate metrics (or combinations of metrics).

VIII. Conclusions

This paper presented preliminary research results towards developing a mathematical formulation for effective autonomous reconstitution of compromised cyber systems. The proposed formulation enables the application of classical optimization approaches to determine optimal choices for reconfiguring and improving the resilience of the cyber-system after one or more natural faults or adversarial attacks.

Preliminary results of simulation studies indicated that the proposed approach may be viable for reconstitution of compromised cyber systems. Unlike existing techniques, the proposed approach is seen to be viable even when faults occur during the reconstitution process. However, the simulation studies to date included several assumptions that will need to be relaxed if the results are to be applicable more generally. The relaxation of these assumptions, and the evaluation of the proposed formulation under these more general conditions, will be the focus of future work.

IX. ACKNOWLEDGEMENTS

This work was supported by the Asymmetric Resilient Cybersystems Initiative at the Pacific Northwest National Laboratory (PNNL). PNNL is a multi-program national laboratory operated by Battelle Memorial Institute for the US Department of Energy under DE-AC06-76RLO 1830.

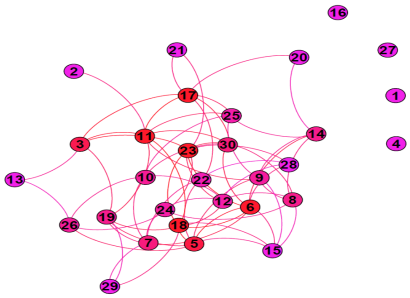

Figure 3. Rendering of a 30-node network with a fault at time-step zero resulting in the isolation of three nodes.

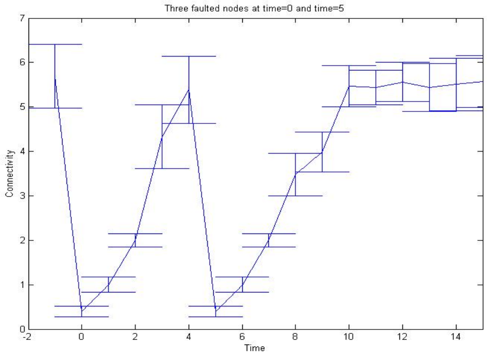

Figure 4. Algebraic connectivity during recovery of the 30-node network. The horizontal axis represents time steps (arbitrary units) after a fault at Time 0. The vertical axis shows the algebraic connectivity. The vertical bars represent the standard deviation over 10 random networks.

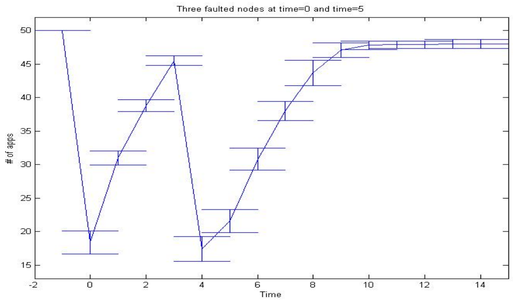

Figure 5. Number of active services during recovery of the 30-node networks

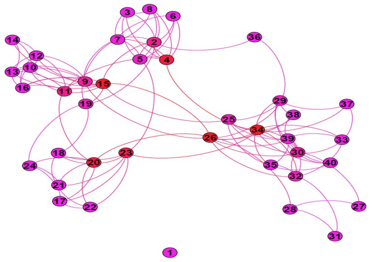

Figure 6. Rendering of the 40-node network. This is an interconnected system of four densely connected networks.

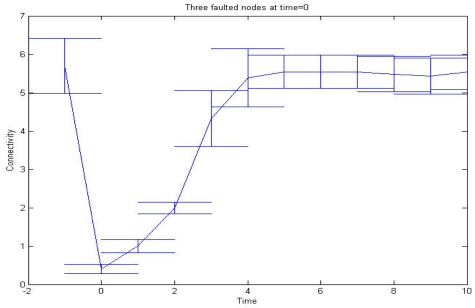

Figure 7. Algebraic connectivity during the recovery of the 40-node network.

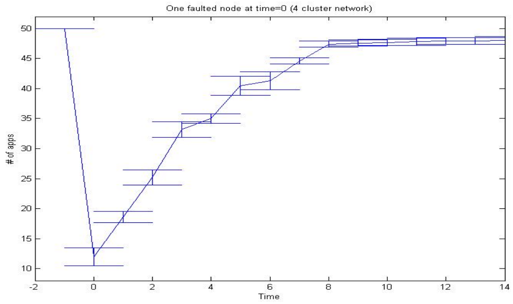

Figure 8. Number of active services during the recovery of the 40-node network.

REFERENCES

[1] S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang, Eds., Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats (Advances in information security 54). Springer New York, 2011.

[2] J. H. Kahan, A. C. Allen, and J. K. George, "An operational framework for resilience," JHSEM, vol. 6, 2009.

[3] G. H. Baker, C. T. C. Mo, and J. G. Voeller, "Time-domain probabilistic risk assessment method for interdependent infrastructure failure and recovery modeling," in Wiley Handbook of Science and Technology for Homeland Security: John Wiley & Sons, Inc., 2008.

[4] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," ACM Trans. Program. Lang. Syst., vol. 4, pp. 382-401, 1982.

[5] P. Beraud, A. Cruz, S. Hassell, and S. Meadows, "Using cyber maneuver to improve network resiliency," in Proc. MILCOM 2011, 2011, pp. 1121-1126.

[6] M. Castro and B. Liskov, "Practical Byzantine fault tolerance and proactive recovery," ACM Trans. Comput. Syst., vol. 20, pp. 398-461, 2002.

[7] G. Bracha, "Asynchronous Byzantine agreement protocols," Inform. Comput., vol. 75, pp. 130-143, 1987.

[8] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," J. ACM, vol. 27, pp. 228-234, 1980.

[9] D. Dolev and H. Strong, "Authenticated algorithms for Byzantine agreement," SIAM J. Comput., vol. 12, pp. 656-666, 1983.

[10] S. Dubois, T. Masuzawa, and S. Tixeuil, "Bounding the impact of unbounded attacks in stabilization," IEEE Trans. Parallel Distr. Syst., vol. 23, pp. 460-466, 2012.

[11] M. Ben-Or, D. Dolev, and E. N. Hoch, "Fast self-stabilizing Byzantine tolerant digital clock synchronization," in Proc. 27th ACM Symp. on Principles of Distributed Computing, Toronto, Canada, pp. 385-394, 2008.

[12] S. Amitanand, I. Sanketh, K. Srinathant, V. Vinod, and C. P. Rangan, "Distributed consensus in the presence of sectional faults," in Proc. 22nd Annual Symp. on Principles of Distributed Computing, Boston, MA, pp. 202-210, 2003.

[13] B. Awerbuch, D. Holmer, and H. Rubens, "Swarm intelligence routing resilient to Byzantine adversaries," in 2004 Int'l Zurich Seminar on Communications: Access - Transmission - Networking, Zurich, Switzerland, February 18-20, 2004, pp. 160-163.

[14] S. Roy, et al., "A survey of game theory as applied to network security," in Proc. 43rd Hawaii Int'l Conf. on System Sciences, HICSS-43, Kauai, HI, January 5-8, 2010, pp. 1-10.

[15] H. Goldman, R. McQuaid, and J. Picciotto, "Cyber resilience for mission assurance," in 2011 IEEE Int'l Conf. on Technologies for Homeland Security, HST 2011, Waltham, MA, November 15-17, 2011, pp. 236-241.

[16] P. Beraud, A. Cruz, S. Hassell, and S. Meadows, "Using cyber maneuver to improve network resiliency," in Proc. MILCOM 2010, Baltimore, MD, November 7-10, 2011, pp. 1121-1126.

[17] A. V. Oppenheim and R. W. Schafer, Discrete-time Signal Processing, 3rd ed.: Prentice-Hall, 2010.

[18] D. Bodeau, et al., "Cyber resiliency metrics, version 1.0, rev. 1," The MITRE Corp, Bedford, MA, MP120053, Rev. 1, 2012.

[19] M. Crouse and E. W. Fulp, "A moving target environment for computer configurations using genetic algorithms," in 4th Symp. on Configuration Analytics and Automation (SAFECONFIG), Arlington, VA, October 31-November 1, 2011, pp. 1-7. Article 6111663.

[20] S. Gilbert, J. Saia, V. King, and M. Young, "Resource-competitive analysis: A new perspective on attack-resistant distributed computing," in Proc. 8th Int'l Workshop on Foundations of Mobile Computing, Madeira, Portugal, July 19, 2012, pp. 1-6.

[21] A. Antoniou and W.-S. Lu, Practical Optimization. New York: Springer, 2007.

[22] V. V. Vazirani, Approximation Algorithms: Springer, 2004.

[23] A. Jamakovic and S. Uhlig, "On the relationship between the algebraic connectivity and graph's robustness to node and link failures," in 3rd EuroNGI Conference on Next Generation Internet Networks, Trondheim, May 21-23, 2007, pp. 96-102.