# Organic Resilience for Tactical Environments

Marco Carvalho[1], Tom Lamkin[2], and Carlos Perez[1]

[1] Institute for Human and Machine Cognition, Ocala, FL
{macarvalho,cperez}@ihmc.us
[2] Air Force Research Laboratory, Dayton, OH
thomas.lamkin@wpafb.af.mil

**Abstract.** In this paper we present a tactical defense infrastructure for mission survivability based on three core inspirations from biological systems: multi-potentiation, feedback mechanisms, and redundancy. In tactical operational environments, these concepts may be realized through a combination of capabilities that include (1) dynamic allocation of resources for mission execution, (2) detection and identification of attacks and their effects, and (3) information sharing for system adaptation. As a proof-of-concept we introduce an extensible, multi-layer defense infrastructure inspired in the self-organization and resilience properties of biological systems. Two defense strategies are considered to validate the proposed model: a fast response consisting on rebooting a compromised system from a reference system image; and a slower response involving a process of identification of the attack, which then allows the node to change its base configuration and reboot to a state that is potentially immune to the same attack. Our experimental results show that the second strategy improves the overall resilience of the system for ongoing attacks after an initial exposure phase.

**Keywords:** Organic Computing, Biologically-Inspired Resilience, Tactical Networks, Resilient Systems, Distributed Control.

## 1 Introduction

One of the hallmark characteristics of biological systems is their capacity to adapt and evolve to environmental changes and competing pressures from peers and adversaries. Self-organization and adaption happens at all levels of scale and complexity, from intra-cellular signaling through larger scale biological systems and social interactions between individuals and groups.

In order to adapt, systems must be able to interact with the environment and respond to localized damages or perceived threats. In a general sense, this goal can only be accomplished if the system is capable to sustain some level of damage, allowing it to perceive, identify and adapt to the problem. These are capabilities that enable biological systems's intrinsic resilience to external attacks and environmental conditions. Similarly, in our view, mission critical infrastructures should be able to seamlessly absorb localized attacks with minimum impact to the ongoing tasks, while isolating and responding appropriately.

J. Suzuki and T. Nakano (Eds.): BIONETICS 2010, LNICST 87, pp. 22–29, 2012.

In this paper we introduce a multi-layer defense infrastructure that realize these capabilities for a mobile tactical environment.

## 2 An Organic Computing Approach to Mission Survivability

The concept of organic computing [1] has been recently introduced to represent and control a computational systems as a living organisms, capable to self-organize, regulate and adapt in order to survive and achieve its goals.

For tactical operational settings, mission survivability is often defined as the ability to maintain the execution of the mission, even under unexpected adverse conditions, localized failures, or attacks. A system capable to identify local failures and block external attacks in a timely manner is thought to be capable to maintain the execution of its mission-critical applications to their successful and timely completion.

The problem is further complicated in tactical environments, where the lack of a fixed coordination infrastructure, the tenuous definition of system boundaries, and the dynamic nature of the computation and communication infrastructures make it very difficult to rely on conditional approaches to improve system resilience. The security requirements for tactical operation environments are significantly different than those normally defined for infrastructure networks [2], requiring new approaches for system defense and mission assurance.

For such environments, the system's defense infrastructure must be as fluid and adaptive as the system itself.An adaptive defense infrastructure for tactical MANETs must be able to learn from successful attacks, by quickly identifying, localizing and isolating the attack to devise a defense at runtime, while ensuring that overall mission requirements continue to be met.

## 3 Mission Survivability in Tactical Environments

In the context of this work, a mission is defined as an ordered set of tasks that must be performed by the computational environment in a timely fashion. While there have been several approaches for mission modeling and representation, including standards such as Business Process Modeling Notation (BPMN) [3], for the sake of simplicity we will choose to describe a mission simply as an array of symbols, where each symbol represents a task and the order defines their interdependencies.

As illustrated in Figure 1, node 1 is tasked with the mission described by the sequence AFG, where each of the symbols represent a self-contained task for the mission. For a tactical service oriented architecture (SOA), the orchestration of services A, F and G must be carried out on the fly, with minimum cost and coordination overhead, and maximum performance possible. In Figure 1, three nodes collaborate to jointly execute the mission, each of them handling one of the sub-tasks A, F and G.

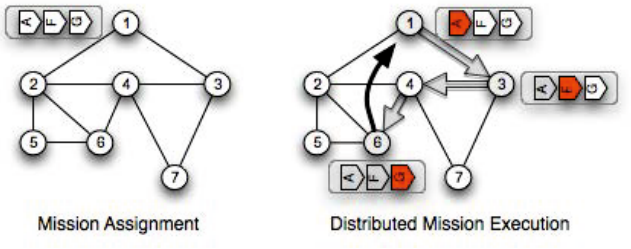


**Fig. 1.** Distributed mission execution through runtime distribution of tasks

The allocation of resources for task execution is handled by resource coordination and orchestration services. In tactical network environments, the allocation of services must account for the unreliability of links and topology dynamics [4]. The defense of such infrastructures, however, is not a trivial task. There are multiple attacks at the network [5] , access control [6], physical [7] and also at the host level. In this environment, nodes must be able to detect, isolate and identify a problem that may affect the performance of the mission. These tasks must be carried while maintaining mission execution with minimal disruption.

## 4 Defense Infrastructure for Mission Survivability

Our proposed defense infrastructure is composed of three components (Figure 2). The first component manages the dynamic allocation of resources for mission execution. The second component is responsible for detection and identification of the potential attack. The third component coordinates the sharing of information about the attack, ensuring that a collective response (if appropriate) can be enforced, and that nodes that are functionally similar to the victim can be reconfigured to prevent a similar attack. A collective response to an attack may include, for instance, modifications in routing weights to disfavor the use of nodes that may have been compromised.

While simultaneously supported and coordinated, the proposed defense infrastructure must be loosely couple to prevent a cascade failure in the even that one of the components becomes temporarily impaired or permanently compromised. As conceived, the coordinated operation of all three components is necessary to enable a comprehensive response and system resilience, each component will also operate independently with limited performance gains, ensuring a graceful degradation of the survivability infrastructure itself.

Automatic resource and service re-allocation in response to localized failures is common practice in Grid environments, and has also been previously proposed for enterprise [8] and tactical [9] environments. However, in general, a change in allocation strategy happens only when degradation (or failure) has taken place and the impact on the mission has been noted, there's generally no predictive re-allocation based on increased risk of an attack or failure, learned at runtime from novel attacks. Our proposed approach leverages and extends such dynamic

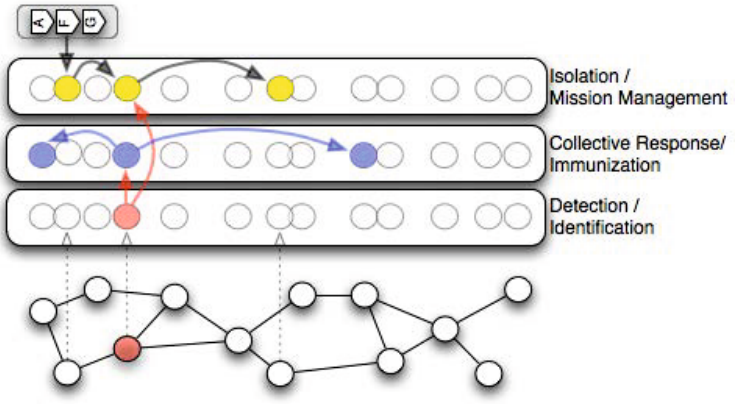


**Fig. 2.** Multi-layer approach to organic resilience

allocation strategies in the literature to enable proactive reallocation based on online risk estimation.

In the general case, attacks are detected indirectly, through their effects on the mission (effects-based detection),but there are related research efforts on mission mapping [10] that can also provide a better assessment of the impact of localized failures to the overall mission by matching the mission workflow with resource requirements and availability.

There are then two possible defense strategies for a node that has detected local damage. The first strategy takes significantly less time, and consists on rebooting from a reference system image, but this strategy yields no details about the attack. The node reboots just as vulnerable to the attack as it was before the detection. The second strategy involves a process of identification of the attack, which then allows the node to change its base configuration and reboot to a state that is potentially immune to the same attack.

## 5 Experimentation Setup and Simulation Environment

The proposed approach was implemented and evaluated in NS3. The simulations were executed over synthetic scenarios with random topologies and random mobility models and also over scenarios based on a dataset collected from a live exercise at the US Army National Training Center (NTC). The NTC dataset describes a 24-hour long exercise including approximately 480 nodes between friendlies and enemies. In addition to providing a realistic operational scenario, the NTC dataset includes the mobility information of all participating nodes, and has been used in other research efforts for mobility modeling.

Each scenario was executed with 20 different seeds and the results were averaged out across those seeds. The metric of interest for comparing results on similar scenarios is the percentage of completed missions at any given moment of

the simulation. First a baseline was computed by running all the scenarios with no attacks during the whole simulation. This baseline metric provides the upper bound for the operational limits of the system, and it is labeled as *Baseline* in the result charts. It illustrates the performance of the system (in terms of number of completed missions) under normal operations. A second baseline metric (labeled as *Baseline No Reboot*) provides the lower bound for the operational limits of the system.

In addition to a short term response (second baseline), an adaptation strategy can be used to enhance the resilience of the system to subsequent attacks. One strategy is to randomize the configuration of re-instantiated services and nodes. A second approach is to provide an immunization capability that will drive the mutations of re-instantiated servers to become resistant to previous attacks. In our experiments, we have opted for an immunization strategy.

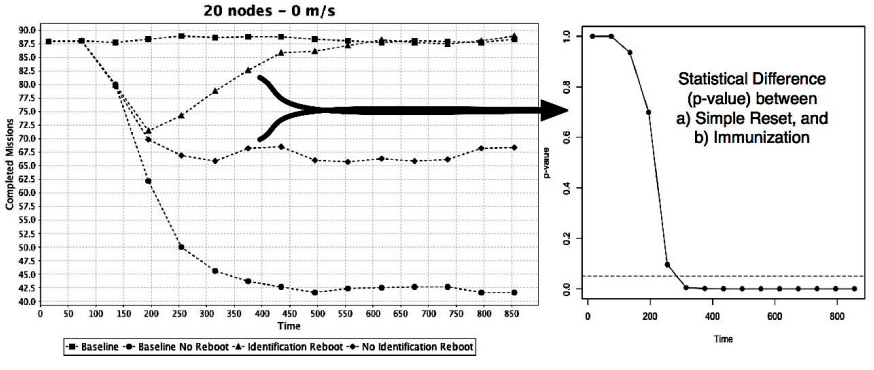


**Fig. 3.** Results for static scenario of 20 nodes

Figure 3 shows the results for the static scenario (no mobility). The results series labeled as *Identification Reboot* and *No Identification Reboot* represent the results for the strategies with and without immunization, respectively. In the first case, nodes detect and identify the attack and then reboot from a previously known safe state in order to maintain mission requirements. The attack detection happens indirectly (through the effects of the attack) and the identification happens by correlating the detection with the current state of the node (represented as a short string). The process takes some time, during which the services provided by the node are degraded.

In the second case (immunization), the node identifies a "mutation" strategy that is likely to make it less vulnerable to the same attack. For our simulations, the state of a node is represented by its 4-bit DNA sequence and defines how vulnerable a node is to a given attack. The immunization process involves having a node that has been attacked to announce its current DNA to its peers, which will drive "similar" nodes to mutate in order to become resistant to the attack.

A 4-bit DNA in this example represents, for instance different OS and software combinations (Web Server and libraries) of a functional node.

In the scenario illustrate in Figure 3, the strategy involving immunization starts with results close to the non-immunization strategy but then it improves until getting close to the upper operational boundaries of the system. Figure 3 shows how the $p$-value changes across time for a $t$-test of difference in percentage of completed missions for the immunization and no-immunization strategies. The dotted line in the chart represents the $p$-value of 0.05, for a confidence interval of 95%. So, approximately 300 seconds into the simulations (around 100 seconds after the attacks start), the difference in performance between the two defense strategies becomes statistically significant.

Figure 4 shows a different network scenario created from a subset of the National Training Center (NTC) dataset consisting of the last 2 hours of two relatively small areas with 28 nodes. The 2 hours of the scenario were compressed into 1200 seconds, to have the same number of missions as for the synthetic scenarios.

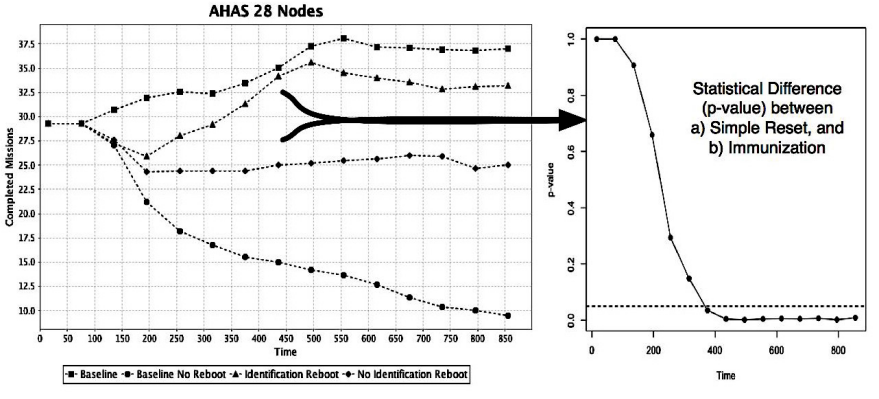


**Fig. 4.** Results for NTC scenario of 28 nodes

Out of the total number of nodes in the scenario, 12 nodes are selected as attackers, and 10 of the remaining nodes are issuing missions to be executed roughly every 2 seconds until completing 400 missions each. The missions and the attacks have the same characteristics as for the synthetic scenarios. Like int he previous case, the NTC scenarios were also executed with 20 different seeds and the results were averaged out across those seeds. In this scenario the system performance starts improving, but then it shows a small decrease. This is due to the fact that the topology starts with a few isolated clusters and then, due to mobility some of these clusters get connected, providing temporarily more resources to execute missions, but later a few get disconnected again, degrading the performance again.

In general, as illustrated in our preliminary results, the immunization strategy provides a gain after some initial period of time in comparison with a simply reactive approach. The results depended on the time required for the identification of the attack and definition of the adaptation strategy. For our tests we estimated an arbitrary relationship of one to four. Meaning that if the time required for detecting the damage from an attack is one second, the time required to identify the attack for immunization is four times longer. In our results, the simple response strategy *No Identification Reboot* requires only the detection, while the immunization strategy *Identification Reboot* requires the identification of the event. The objective in this example is defined as the overall number of missions executed by the system and since each mission has a time to live, it implicitly includes the time requirements for mission execution.

It is important to note that the balance between the time required to learn and disseminate an immunization strategy, the time required for damage detection and the rate of attacks will play an important role on defining the appropriate strategy (or mixture of strategies) in a more general case. While we envision simple approaches for choosing different strategies for different operational conditions, we have not yet fully evaluated these effects, which will be part of our continued work in this topic.

## 6 Conclusion and Future Work

In this work we presented a defense infrastructure that loosely couples a fast-response mission management capability with a threat isolation, identification and mitigation capability operating, in coordination, at a slower pace.

While the infrastructure has been designed for improved mission survivability, it is important to highlight that our approach is vulnerable to other attack strategies, which must be managed by complementary capabilities.

As future work, we would like to explore other mission allocation approaches that provide better online risk estimation. For example, other approaches using distributed reinforcement learning have been used with promising results [11].

It would be also of interest to investigate the use of distributed anomaly detection strategies for isolating malicious nodes and finding the root causes of the attacks in a coordinated fashion. This type of strategies might help mitigate some vulnerabilities of the infrastructure, like detecting a slow overall degradation of the infrastructure.

## References

1. Muller-Schloer, C.: "Organic computing" on the feasibility of controlled emergence. In: CODES+ISSS 2004: Proceedings of the International Conference on Hardware/Software Codesign and System Synthesis, pp. 2–5. IEEE Computer Society, Washington, DC (2004)
2. Carvalho, M.: Security in mobile ad hoc networks. IEEE Security and Privacy 6(2), 72–75 (2008)

3. Wong, P.Y., Gibbons, J.: A process semantics for BPMN. Draft (2007)
4. Carvalho, M., Winkler, R., Perez, C., Kovach, J., Choy, S.: A cross-layer predictive routing protocol for mobile ad hoc networks. In: Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ser. Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, vol. 6981 (May 2008)
5. Agrawal, P., Ghosh, R.K., Das, S.K.: Cooperative black and gray hole attacks in mobile ad hoc networks. In: ICUIMC 2008: Proceedings of the 2nd International Conference on Ubiquitous Information Management and Communication, pp. 310–314. ACM, New York (2008)
6. Gupta, V., Krishnamurthy, S., Faloutsos, M.: Denial of service attacks at the mac layer in wireless ad hoc networks (2002)
7. Zhou, Y., Wu, D., Nettles, S.M.: On mac-layer denial of service attacks in ieee 802.11 ad hoc networks: analysis and counter measures. Int. J. Wire. Mob. Comput. 1(3/4), 268–275 (2006)
8. Lardieri, P., Balasubramanian, J., Schmidt, D.C., Thaker, G., Gokhale, A., Damiano, T.: A multi-layered resource management framework for dynamic resource management in enterprise dre systems. J. Syst. Softw. 80(7), 984–996 (2007)
9. Carvalho, M., Pěchouček, M., Suri, N.: A Mobile Agent-Based Middleware for Opportunistic Resource Allocation and Communications. In: Thompson, S.G., Ghanea-Hercock, R. (eds.) DAMAS 2005. LNCS (LNAI), vol. 3890, pp. 121–134. Springer, Heidelberg (2006)
10. Musman, S., Temin, A., Tanner, M., Fox, D., Pridemore, B.: Evaluating the impact of cyber attacks on missions. In: 5th International Conference on Information Warfare and Security, April 8-9, pp. 446–456. Air Force Institute of Technology, Wright-Patterson AFB (2010)
11. Carvalho, M.: A distributed reinforcement learning approach to mission survivability in tactical manets. In: CSIIRW 2009: Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research, pp. 1–4. ACM, New York (2009)