A Framework to Evaluate Cyber Agility

Lisa M. Marvel[1], Scott Brown[2], Iulian Neamtiu[3], Richard Harang[1], David Harman[1], Brian Henz[1]

[1]U.S. Army Research Laboratory
APG, MD 21005
lisa.m.marvel.civ@mail.mil

[2]Secure Mission Solutions
North Charleston, SC 29406

[3]Univ. of California, Riverside
Riverside, CA
neamtiu@cs.ucr.edu

Abstract— **In this paper, we propose a framework to help evaluate the cost and utility of cyber agility maneuvers within networks that have constrained resources such as bandwidth and energy (e.g., MANETs). Many new methods of cyber agility and defensive maneuvers have been proposed; however, a framework to evaluate cost and utility of these maneuvers in the context of mission success is lacking. We outline preliminary considerations such as mission goals, operating conditions and maneuvers to be evaluated. Then, we introduce notional measures of health, security and capability and their interrelationship resulting in an initial framework design. A simple defensive cyber operation of securing a critical communication path for some time duration, both with and without the presence of a detected infection, is provided to illustrate the framework components and mission considerations that must be made before selecting a sequence of maneuvers.**

Keywords—Cyber Security, Computer Network Defense, Computer Network Operations

I. INTRODUCTION AND BACKGROUND

In addition to its original role as a communications medium, the cyber domain is rapidly expanding to become both a force multiplier and a field of operations in its own right. As communications and networked components fuse into this new domain, the security of these networks becomes increasingly vital to the success of the mission and the security of our troops.

Over the past few years, there have been multiple efforts to develop methods for moving target defense [1] with the first moving target defense workshop held in 2014 [2] among others. The security and science of Agility [3] outlined considerations for cyber agility and how moving target methods can be used to accomplish cyber missions.

Cyber agility maneuvers can be used to help lead mission success and minimize the impact of adversary actions. The overall effectiveness of cyber operations can be enhanced with the help of introducing cyber agility maneuvers both reactive and proactive. Cyber maneuvers, such as moving target defense, have been explored to help increase the security of a network, but a method to evaluate their utility and aid in the selection a sequence of maneuvers does not exist. Ideally, methodologies that guide the selection of maneuvers in the dynamically changing cyber environment in which priorities change in response to current mission requirements are needed.

Our research attempts to address a few open questions. Namely, how can we measure the utility (cost/benefit) of cyber maneuver techniques? Can we define composite measures to help guide agility decisions (and sequences of agility decisions) for a system that consists as a collection of nodes in a network? How do we select the best maneuvers at time *t* for node *n* in the context of a local or global mission success? In addition, how do we inform our agility maneuver selection so that we consider costs to constrained nodes and networks? What are the additional considerations when energy and link throughput may vary?

The remainder of the paper is organized as follows. We first present preliminary considerations such as mission goals, operating conditions and maneuvers to be evaluated. Then, we introduce notional measures of health, security and capability and their interrelationship for the maneuvers. A brief discussion on optional maneuver selection follows along with suggested methods for calculated cost. We then use the framework to evaluate the potential maneuvers for two scenarios, namely vulnerability only and when vulnerability has been exploited and the infection is detected. Finally, we conclude with a summary of our findings and propose future work.

II. FRAMEWORK BASICS

In this section, we begin by defining a simple defensive cyber mission that consists of primary and secondary goals. The initial set of agility maneuvers from which we can choose are also defined, as well as the operating conditions for the mission.

A. Primary Goal

The primary goal for this mission is to secure a critical communication path through a network for some time duration to transfer vital information. For example, this could be in preparation for a critical file transfer between two nodes in the network requiring traversal of intermediate nodes. While we are securing this critical path, we have the option of selecting agility maneuvers that will maximize the capability of nodes on critical path while minimizing energy consumption expended to perform the maneuvers in a resource-constrained environment.

B. Secondary Goal

The secondary goal is to secure the entire network in minimal time while maximizing capability of network nodes and minimizing energy consumption.

C. Operating Conditions

For the mission, we consider two operating conditions. The mission occurs:

1) in the presence of a known vulnerability for which a patch is present within the network

2) in the presence of a detected infection that propagates through the network exploiting a known vulnerability for which a patch exists and is present within the network

D. Potential Manuevers

To demonstrate the cost/utility evaluation within the framework, a set of agility maneuver options applicable to the mission are selected. They are as follows:

1) p = Patching: A software patch is delivered to a node and applied to oleviate a vulnerability

2) h = Software Healing: the vulnerability is mitigated by rewriting the software on the node itself

3) q = Function Quarantine: The vulnerability is mitigate by rewiting the software on the node to restrict code execution by isolated in the function containing vulnerable within the software

4) b = Node Blocking: The node is blocked and considered unreachable by other nodes in the network.

III. Notional Metrics and Relationships

In this section, we introduce notional measures of Health, Capability and Security. As the nodes and networks change, the notional measure will change in response. We define initial relationship for each of these notional measures as a function of the agility maneuver set as well. These are be used as a basis for evaluation in the framework.

A. Notional Health Measure

Each node is assigned a Health state, H. The initialized state for all the nodes is immune/patched. As a new vulnerability (or an infection that exploits an existing vulnerability) is discovered, a nodes health changes accordingly:

1) p = Patched/Immune: The node is running software that has been patched and is immune to the vulnerability/infection.

2) v = Vulnerable: The node is running software that contains a known vulnerability that may be exploited.

3) i = Infected: The node is running software that contains a known vulnerability and malware has been delivered to the node.

4) s = Susceptible: The node is running software that contains a known vulnerability that may be exploited, and they are in direct contact with a known infected node.

The relationship among the various Health states, $H(\cdot)$, is as follows:

$$H(p) \geq H(v) \geq H(s) \geq H(i)$$

A diagram of the Health state changes is shown in Figure 1.

B. Notional Capability Measure

In addition, each node is assigned a Capability state, $C(\cdot)$. The initialized Capability state for all the nodes is immune/patched. Capability is a function of the agility maneuver selected for the node and the patched/immune and infected Health states. The relationship is as follows:

$$C(p) \geq C(h) \geq C(q) \geq C(i) \geq C(b)$$

In this relationship, we are assuming that the best node Capability results from the patching/immune Health state.

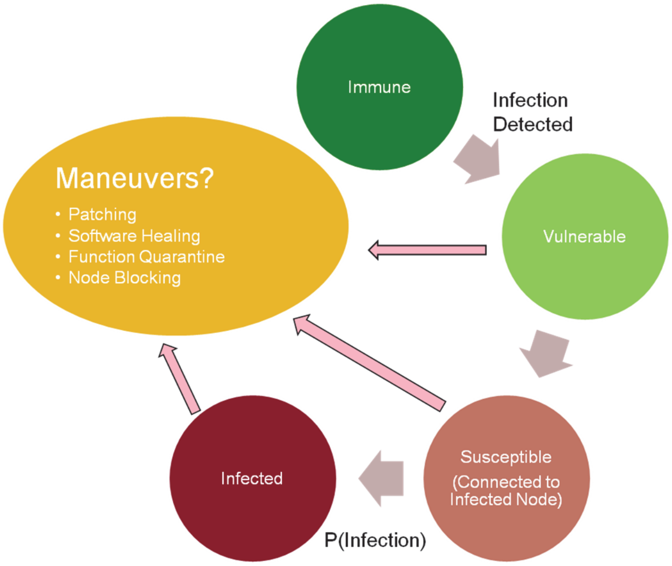

Figure 1 Health State Transitions

The next highest capability results from software healing, which can rewrite the software in a way that the vulnerability is no longer exploitable while the functionality of the software is still available. This is followed by the Function Quarantining option, which rewrites the software making the vulnerability function unusable by the node. This will reduce the capability of the node because the functionality provided by the vulnerable function is no longer available. An infected node has a reduced Capability, although it may still contribute to the mission, e.g., by forwarding packets. Capability is further reduced if the node is blocked from other nodes in the network, preventing it from contributing to the mission entirely.

C. Notional Security Measure

Lastly, we present a notional measure of security, $S(\cdot)$. We can reason that the security level of each node contributes to the overall security of the network. The measure is also a function of the agility maneuver selected for the node and the patched/immune and infected Health states following this relationship:

$$S(p) \geq S(b) \geq S(q) \geq S(h) \geq S(i)$$

Here, a patched immune system provides the greatest security measure of the potential maneuvers, followed by the blocking node, which does not negatively impact the security of the network if it is not connected. The node Security measure may decline with Function Quarantine, Healing (due to the on node modification of software necessary for those maneuvers) and lastly the infected node.

Note: The Health, Capability and Security measures are notional, and although we have reasoned about their relationships, they are presented for example. These relationships could change as maneuvers are introduced and refined and the mission goals and operating condition are modified.

D. Composite Notionals Measures

State changes are performed in steps that can have varying duration. We can evaluate the notional measures at each step. For instance, let variable $\mathcal{H}_s$ denote the composite health of the network at step $s = \{0, 1, 2, \dots, w\}$, where *w* represents the maximum number of steps for the network to achieve full security and capability for all possible sequences of agility maneuvers. Let *N* be a set of nodes in a network of size *n* with nodes $j = 1, 2, \dots, n$; at each step, *s*, the notional measures can be computed. For instance, if $H_{j,s}(\cdot)$ is the health current state of node *j* at step *s*, then the health state of the network is expressed as the average by accumulating all the states health for all nodes in the networks:

$$\mathcal{H}_s = \frac{1}{n}\sum_{j \in N} H_{j,s}(x) \qquad (1)$$

where $x \in \{p, v, s, i\}$. For the critical path within the network, we define $K \subset N$ as the set of nodes on the critical path; the notional health of the critical path at step *s* would be as follows:

$$\mathcal{H}_{K,s} = \frac{1}{|K|}\sum_{k \in K} H_{k,s}(x) \qquad (2)$$

where $|K|$ represents the cardinality of the set *K*.

In addition, the notional health measure for the network can be established by accumulated all the states as follows, where $x \in \{p, v, s, i\}$:

$$\mathcal{H}_N = \frac{1}{nw}\sum_{j \in N}\sum_{s=1}^{w} H_{j,s}\,(x) \qquad (3)$$

Correspondingly, for a selection of maneuvers and starting/ending states represented as $y \in \{p, h, q, b, i\}$, the notional measure for capability (and security) can be expressed as follows:

$$\mathcal{C}_s = \frac{1}{n}\sum_{j \in N} C_{j,s}(y) \qquad (4)$$

$$\mathcal{C}_{K,s} = \frac{1}{|K|}\sum_{k \in K} C_{k,s}(y) \qquad (5)$$

These measures can be extended for network as a whole as in (3) over the domain *y*.

IV. Optimal Selection of Manuevers

Ultimately, we seek the selection of the sequence of agility maneuvers for each node in the network, which will result in maximizing health, security and capability of the network while minimizing the cost in terms of energy and time.

As an example, if we optimize each node independently and assume a mission prioritizes security over capability while minimizing energy consumption, we first find the set of maneuvers that will maximize capability:

$$M' = \{M: C(M) = \max C(z)\} \qquad (6)$$

Then we find the subset of these maneuvers that will maximize security:

$$M'' = \{M: S(M) = \max_{z \in M'} S(z)\} \qquad (7)$$

Lastly, we find the subset of these maneuvers that minimize energy consumption:

$$M_j = \arg\min_{M \in M''} P(M) \qquad (8)$$

Following this logic, we define a function γ that represents the tradeoff between energy, capability and security as determined by the mission. We can optimize to find the best set of maneuvers: Let $\mathbb{M}_j$ represent the ordered set of all possible agility maneuvers taken by node $j \in N$, then we can express this as some function γ, where

$$\gamma_t : P, C, S \rightarrow \Re \qquad (9)$$

$$\arg\max_{M_j \in \mathbb{M}_j} \gamma_t(M_j, t) \qquad (10)$$

Here, the subscript *t* on γ indicates that we may alter our cost/utility function as time goes on and the mission develops as will occur during dynamic cyber missions. Consider the overly simplistic example, where at the start of a mission, security may be prioritized over capability as in

$$\gamma_0 = -\mathrm{P} + 1\mathrm{C} + 10\mathrm{S}. \qquad (11)$$

At a critical point in the mission, a commander may elect to sacrifice energy and security constraints in the short term to ensure the provision of critical capability. This can be reflected as follows:

$$\gamma_{t'} = 0\mathrm{P} + 10\mathrm{C} + 1\mathrm{S} \qquad (12)$$

indicating that energy consumption will not be penalized while capability is prioritized heavily over security. This is in concert with the dynamic end states described in [3].

V. Calculating Costs

If the dynamic nature of cyber is ignored, then offline costs can be combined with selected values for standard energy consumption, network link quality, infection propagation, and communication necessary for the mission essential tasks. Then, selection of best maneuvers can be pursued using a dynamic programming approach. However, if we want to truly capture the dynamic nature of the environment costs (a combination of offline and online costs) required, this must be calculated using dynamic interactive models of battery usage, communication, link quality, etc.

For constrained environments, we can express costs of all maneuvers in terms of energy in Joules to align with battery capacity, which is Watt hours (Wh). Using energy consumption, our cost measure incorporates time insofar as communication transmission; so, in many cases, when we minimize energy we also minimize time duration.

A. Calculated Some Costs Offline

It is reasonable to assume that energy estimates to heal vulnerable software, quarantine a function and update routing tables to block a node can all be calculated offline using nodes similar to that in the network. In addition, it is reasonable to assume they can be represented as a constant. This is also the case when calculating the energy cost necessary to apply the patch.

B. Calculating the Cost to Patch

The costs required to patch a node in a mobile network can be segmented into two values: the cost to transmit the patch through the network to the node and the cost to apply the patch, p_a, once it is received by the node. To calculate the former, we assume the network consists of identical nodes with varying energy budgets and link qualities. We define the data transfer rate in bytes/second for a link between node i and j as $T_{i,j}$. The energy required to receive data at a node is p_{rx} , and the energy required to transmit data is p_{tx}, both in units of Joules. As an example, assume that the network nodes are linked between Node m and Node r with $r-m$ intermediate hops. Node m holds the patch needed by Node r. Given a patch of size, b bytes, the cost to patch can be calculated as follows:

$$P_P(m,r) = \sum_{i=m}^{r-1} \frac{b}{T_{i,i+1}}(p_{tx} + p_{rx}) + \ p_a \qquad (13)$$

Again, it is reasonable to assume that the energy required for applying a patch, p_a, can be calculated offline, and it can be represented as a constant value.

VI. Experimental Evaluation

Evaluation of agility maneuvers quickly becomes complex and intractable. In particular, this is the case when dealing with varying energy budgets, dynamic mobile networks such as MANETs with varying link qualities. Using the agility evaluation framework and NS-3 simulations experimentation evaluations can be performed.

To evaluate our candidate set of agility maneuvers for our mission from Section II.D, we constructed a 10-node network using NS-3. A diagram of the network is depicted in Figure 2 showing identifying node number, connected links and critical path (yellow links).

Initially, each node in the simulation is identical and contains a battery model, link information and current state. The battery model is used to simulate drawdown of the battery over the course of the simulated mission. Every node begins with an initial energy value of 63 watt hours (Wh). The connections that exist between each node each are represented using a two metrics, delay in milliseconds (ms), and data rate in megabits per second (Mbps). In the experiment, each link contains a delay of 100 ms and a data rate of ~10 Mbps. In addition, when a packet is sent, it must be forwarded by each node that lies between the source and destination node.

We investigate two operation conditions scenarios described in Section II.C in which agility maneuvers are used to maintain/increase capability while securing the network, maintaining/increasing capability and improving health. The first scenario is equivalent to patch management: some vulnerable nodes exist in the network and patched/securing is required. For the second scenario, a vulnerability exists for which there is a patch; however, before nodes in the network can be patched, an actual infection exploiting this vulnerability is detected and can propagate with some probability.

There are several selections made at the start of the simulation, including the start and end locations for the mission critical nodes, the source of the patch and the source of the infection (Scenario 2). The critical path is established as the shortest path between the start and end critical nodes. All of the nodes along this path are considered critical nodes. Securing the critical nodes in the network is the primary objective of the simulation and the network in its entirety is secondary. Battery drain occurs each time a packet is sent or received by one of the nodes. This allows measurement of the impact of various approaches for containing the infection. The simulation ends once all of the nodes in the network return to full health.

For the purposes of the simulations, we set the notional measures of Health, Capability and Security to numeric values as specified in Table 1. The value assignments preserve the relationships described in Section III.

Table 1 Simulation Values for Notional Measures

Values	0.0	25	50	75	100
Health	$H(i)$	$H(s)$	---	$H(v)$	$H(p)$
Capability	$C(b)$	$C(i)$	$C(q)$	$C(h)$	$C(p)$
Security	$S(i)$	$S(h)$	$S(q)$	$S(b)$	$S(p)$

A. First Operating Scenario: Vulnerability Discovered

The first operating scenario is equivalent to performing patch management. The first scenario contains no vulnerability-exploiting malware (infection). However, 50% of the network nodes have a known vulnerability, with the other 50% of the network are immune (perhaps they do not possess the vulnerable software). Nodes 0–4 represent the vulnerable nodes (indicated as blue), and nodes 5–9 represent the immune nodes (indicated as green P). Figure 2 shows that each of the vulnerable nodes is spread out at various locations around the network. The critical path nodes are nodes 4, 7, 6, and 1. This scenario allows us to determine the most efficient way to patch the remaining vulnerable nodes given the network topology and energy budget. There are 230 possible maneuver sequence selections in set $\mathbb{M}_j$. We compute $\mathcal{H}_N, \mathcal{C}_N$ and $\mathcal{S}_N$ to find the best maneuver selections for each notional measure.

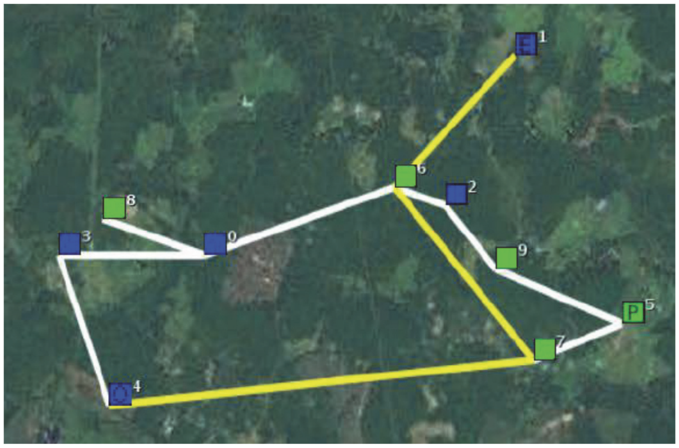

Figure 2 Network Connectivity (Scenario 1)

We can use a heatmap to show how nodes change health state for scenario 1.

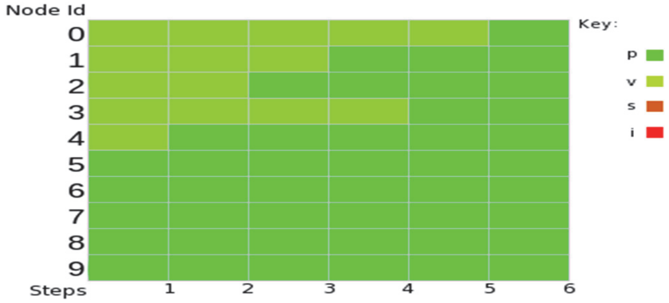

Figure 3 Best Health Heatmap (Scenario 1)

For this scenario, capability and security do not change and only health reflects a vulnerability. The following can be noted in the heatmap: (1) node 4 is selected for patching first because it is on the critical path; (2) then, node 2 is selected, and although is not on the critical path, it is close to node 6; (3) then, node 9, which posses the patch and transmits to node 2, is selected. Node 1, although far from a patch, lies on the critical path and is patched before nodes 0 and 3, which are not critical path nodes. We note that although steps appear to be equal in size, their time duration is variable, and therefore adjacents steps may happen almost simultaneously, whereas others may take longer. It so happens that there is a single manuever selection (run 41 in our data) that produces the best healh, capability and security while minimizing energy expended.

B. Second Operating Scenario: Infection Detected

For the second operating scenario, an infected node is selected along with the probability of infection, denoted *P(infection)*. Nodes communicating with the infected node have a susceptible health state and will in turn become infected with P(infection) = 0.8 for each communication exchange with the infected node. At the start of simulations, attempts are made to minimize the spread of infection by utilizing cyber agility maneuvers. There are 505 possible maneuver sequence selections in set $\mathbb{M}_j$. Again, we compare notional measure for the network to select the best for each measure.

Figure 4 shows the initial state at the start of the simulation. Node 3 is the infected node and nodes 4 and 0 are susceptible. The vulnerability patch is located on node 5. All other nodes in the network are vulnerable to the exploiting malware. The critical path nodes are the same as the previous scenario: nodes 4, 7, 6, and 1. As the malware propagates to the network, nodes continue to change from vulnerable to susceptible and finally infected based on the communication with infected nodes and *P(infection)*. By utilizing cyber agility maneuvers, we attempt to halt or slow the malware infection, secure and maintain capability or first the critical path and then the entire network.

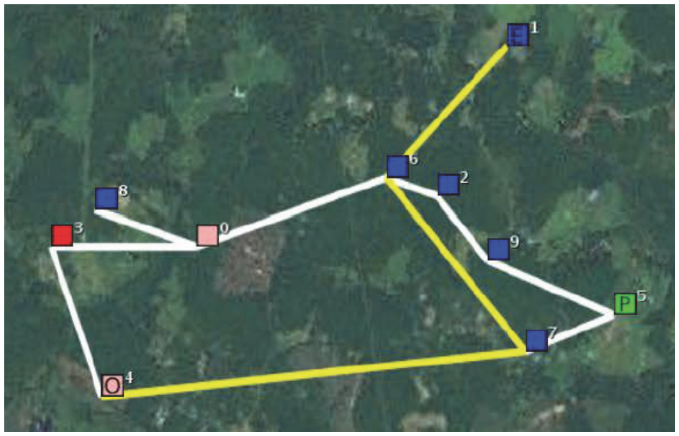

Figure 4 Network Connectivity (Scenario 2)

We can use heatmap (**Figure 5**) to show how nodes change health state for scenario 2 for the best performing manuever selection in respect to health (run 169 in our data).

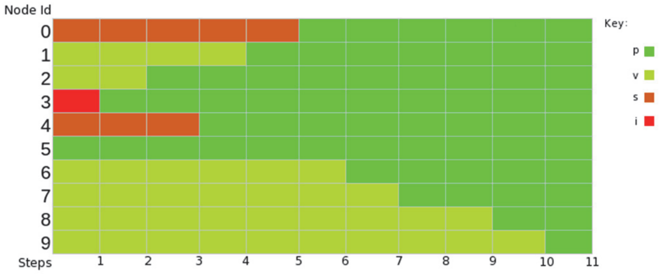

Figure 5 Best Health Heatmap (Scenario 2)

Box and whisker plots are used to show how the best maneuver selection compares with all possible maneuver selections. Figure 6 scenarios are reflected in this chart for health measure.

The heatmap (Figure 7) shows how nodes capability measure changes for the best performing maneuver selection in respect to capability (run 386 in our data).

Lastly, we show a heatmap for the best maneuver selection (run 210) in Figure 9 along with the box and whisker comparison of all maneuvers in Figure 10.

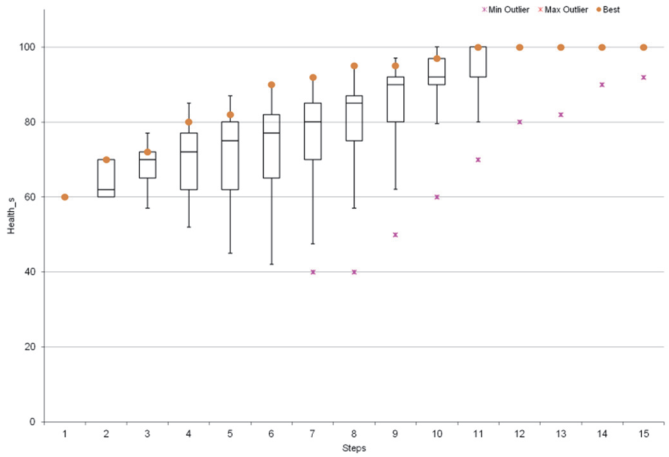

Figure 6 $\mathcal{H}_s$ (Scenario 2)

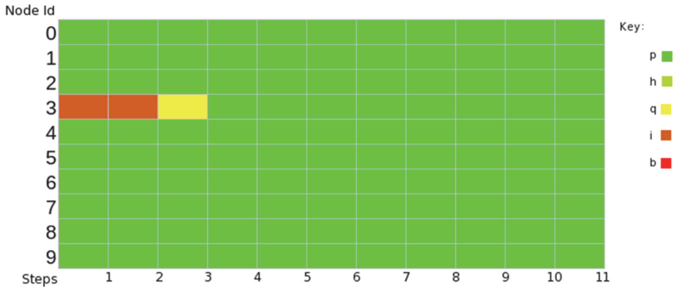

Figure 7 Best Capability Heatmap (Scenario 2)

Box and whisker plots shows the best maneuver selection compares with all possible maneuver selections for capability in Figure 8.

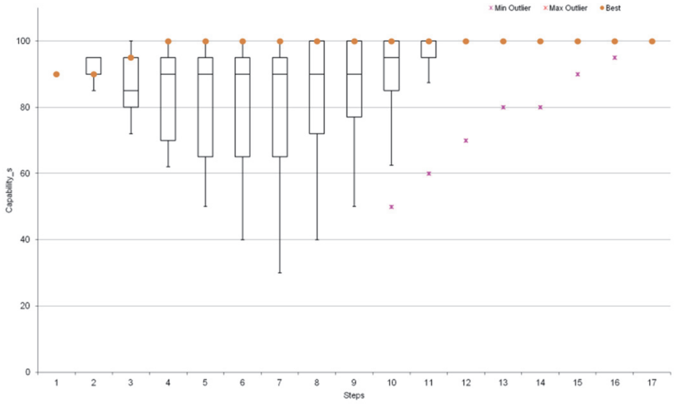

Figure 8 $\mathcal{C}_s$ (Scenario 2)

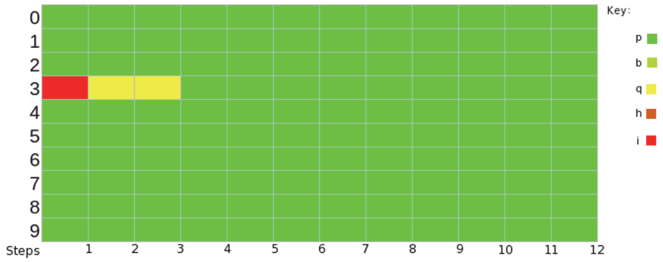

Figure 9 Best Security Heatmap (Scenario 2)

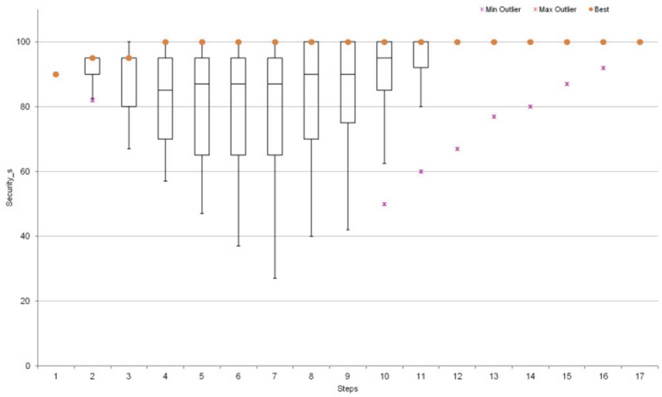

Figure 10 $\mathcal{S}_s$ (Scenario 2)

VII. Conclusion and Future Work

We have presented a framework that can provide a foundation to future work in cyber agility evaluation. We have focused on simple maneuvers and showed how notional measures and their relationship as a function of the various maneuvers can be used to guide maneuver selection. Futhermore, simulations can help to calculate costs in a dynamic network environment where terrain, physical radio communication links, communication volume and routing protocols can be varied.

At this time, we consider a single vulnerability and a single infection in our scenarios. In the future, information about severity of vulnerability/infection and risk could be incorporated into this framework as well as multiple vulnerabilities and infection (with varying propagation rates) as well as multiple mission goals. In addtion, replacement of the notional measure of health, security and capaability with quantifiable metrics should be pursued. For instance, a specification language to define and measure software/node capabilities could help better quanitify a nodes capability.

Acknowledgment

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA) and partially under Contract Number N65235-06-D8847. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] S. Jajodia., A. Ghosh, V. Swarup, C. Wang, X. Wang, (Eds.) Moving Target Defense, Springer Advances in Information Security 2011.

[2] First ACM Workshop on Moving Target Defense (MTD 2014) in conjunction with the 21st ACM Conference on Computer and Communications Security (CCS), Scottsdale, Arizona, November 3, 2014.

[3] P. McDaniel, T. Jaeger, T. F. La Porta, N. Papernot. R. J. Walls, A. Kott, L. Marvel, A. Swami, P. Mohapatra, S. V. Krishnamurthy and I. Neamtiu, "Security and Science of Agility," First ACM Workshop on Moving Target Defense (MTD 2014), Scottsdale, Arizona, November 3, 2014.

[4] T. Azim, I. Neamtiu, and L. Marvel, "Towards Self-healing Smartphone Software via Automated Patching," 29th IEEE/ACM International Conference on Automated Software Engineering, Västerås, Sweden, September 2011