

# MIRA: A Support Infrastructure for Cyber Command and Control Operations

Marco Carvalho, Thomas C. Eskridge  
Harris Institute for Assured Information  
Florida Institute of Technology  
Melbourne, Florida  
mcarvalho,teskridge@fit.edu

Kimberly Ferguson-Walter  
U.S. Department of Defense  
Baltimore, MD  
kjfergu@tycho.ncsc.mil

Capt. Nicholas Paltzer,  
David Myers, and David Last  
U.S. Air Force Research Laboratory  
Rome, N.Y.  
nicholas.paltzer@us.af.mil

**Abstract**—In this paper we introduce MIRA (Mission-aware Infrastructure for Resilient Agents), an extensible and modular infrastructure designed for the support of coordinated cyber operations. MIRA was originally designed to provide the infrastructure for a cyber-command and control and cyber operations. MIRA provides a core set of services that can be extended to allow for operations in specialized or hybrid environments involving different service requirements or operational constraints. The framework was designed to support the easy integration of third party network components, support services, or communication protocols. After a brief motivation about the needs and requirements for command and control in cyber operations, we introduce and discuss the design and implementation of the MIRA framework. The paper follows the discussion with the description of a simple scenario used to illustrate MIRA.

## I. INTRODUCTION

The concept of command and control (C2) is generally associated with the exercise of authority, direction and coordination of assets and capabilities [1]. The concept of C2 encompasses important operational functions such as the establishment of intent, allocation of roles and responsibilities, definition of rules and constraints, and the monitoring and estimation of system state, situation, and progress. More recently, the notion of C2 has been extended beyond military applications to include cyber operation environments and assets.

Unfortunately this evolution has enjoyed faster progress and adoption on the offensive, rather than defensive side of cyber operations. One example is the adoption of advanced peer-to-peer C2 frameworks for the control of malicious botnets and coordinated attacks, which have successfully yielded very effective and resilient infrastructures [2].

Defensive C2 is normally associated with a system's ability to monitor, interpret, reason, and respond to cyber events, often through advanced human-machine interfaces, or through automated actions. Recent research activities in defensive C2 operations are now showing great potential to enable truly resilient cyber defense infrastructures.

In this paper we introduce and describe an agent-based framework called MIRA, designed to support resilient command and control infrastructures for cyber operations. Our focus is on the coordination of dynamic and moving target defenses, enabling C2 infrastructure that must be continually changing and evolving to better respond to an environment, requirements, and adversaries that are also continually changing and evolving [3], [4].

## II. REQUIREMENTS FOR CYBER C2

The goal of Cyber C2 is to increase the resilience of the overall computer network. Resilience is found not in the capabilities of individual defenses, but in the coordinated and contextual use of multiple defenses [5], [6]. We have identified several requirements necessary in order to create resilience in Cyber C2 systems.

**Support Infrastructure for Distributed Communication and Computation.** A key requirement for Cyber C2 is a robust, distributed support infrastructure that enables a separation between the physical structure of resources on the network, and the logical organization of the network and network resources. This can be done in many ways using, for example, custom programming, Service-Oriented Architectures (SOAs), or multi-agent systems [7]. In our approach, illustrated in Figure 1 the support infrastructure for the distributed command and control is provided by MIRA, a multi-agent system designed for resilient cyber operations. In the figure, instances of the MIRA Environment (represented by the containers labeled as E-X) are created across the network, and collectively form the distributed MIRA infrastructure for C2 support. Collectively MIRA Environments will provide resilient core services for the multiple software that will be organized in groups for the monitoring and control of individual sensors and defenses. This abstraction is necessary to support multiple run-time control loops that may be needed to coordinate multiple sensors and defenses.

**Semantic Reasoning.** Reasoning in abstract control loops such as those shown in Figure 1, requires an abstract representation of the network resources, sensors, actuators and defenses, as well as attack vectors, and mission requirements. In MIRA, we represent these resources as OWL (Web Ontology Language) concepts [8], which enables abstract reasoning over network configurations, and specification and enforcement of semantically-aware policies that guide the behavior of the C2 system. Both of these capabilities are needed in order to reason about the effects of adding or removing a defense (cf. [9]) or selecting between equally applicable defenses.

Fully automated control of highly complex systems tends to be brittle against intelligent adversaries [10]. A solution must achieve the right balance between human and automated control. In fact, the best approach is to combine the control methods so that components of both human and automated control systems are operating continually and simultaneously.

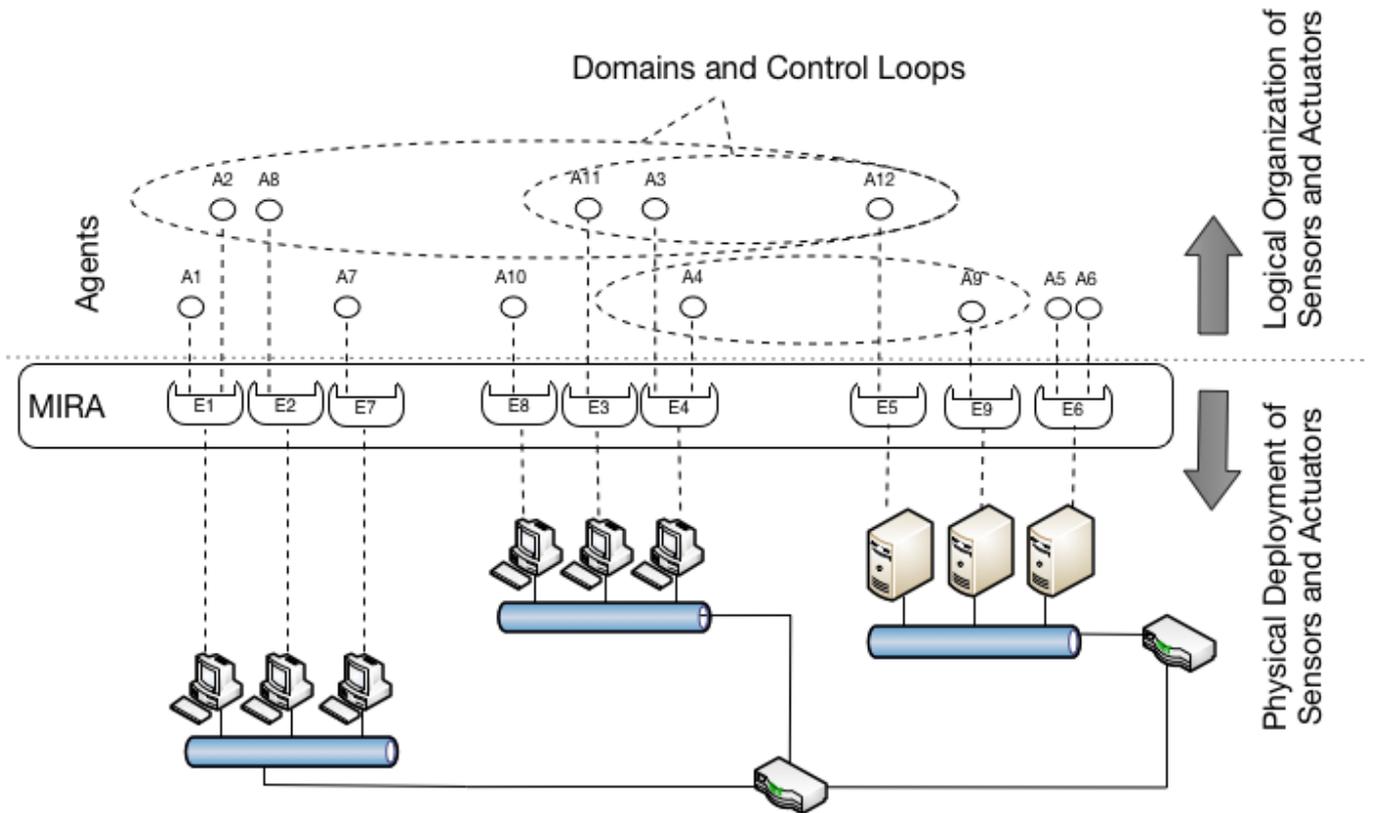


Fig. 1. Logical view of an agent-based command and control framework.

To achieve fast decision rates in critical areas, human input into control decisions is often made using semantic policies.

**Flexible Interface for the Integration of Third Party Sensors and Defenses.** In order to support the requirements for abstract representation and semantic reasoning, it is necessary for a Cyber C2 to be able to integrate new, third party sensors and defenses into the C2 framework. This involves two aspects of integration. First is the integration that allows the C2 system to control the sensor or defense by executing control functions or modifying the parameters of its operation. This requires use of an API, socket, remote procedure call, or file-based interactions. Second is the integration that enables the C2 system to determine when to deploy, enable, and/or modify the operation of the sensor or defense. This involves representing the resource requirements, controls, and parameter ranges and settings for the new sensor or defense in a way that the Cyber C2 can understand, manipulate, and control. In MIRA, this is accomplished by automatically creating an ontological representation when specifying the controls and their ranges. Prior research efforts [4][3] have demonstrated that new sensors and defenses can be integrated with an existing MIRA Cyber C2 system in very little time using this approach.

**Distributed Learning and Coordination Support.** Resilience as defined in [5] requires not only robustness to withstand attacks and the ability to recover back to a pre-attack level of performance, but to reorganize so that the same or similar attacks will not have the same effect on the network and hosts. This requires learning and coordination support so

that an appropriate combination of defenses can be selected and parameterized to ensure maximal performance. MIRA can make use of several algorithms to enable this, including reinforcement learning [11] [12] [13], bayesian model building [14] [15], or hierarchical model building [15][16].

**Advanced Visualization and Human-Computer Interfaces.** Computer network defense is a complicated domain, requiring monitoring and understanding the relationships between hundreds of variables coming from numerous sensors and performance monitors placed throughout the network. It is critical that this information is presented in a way that is not only clear and understandable, but in a way that actually amplifies operator performance [17]. MIRA has several interfaces that show the organization and operation of the MIRA infrastructure [18] and of the traffic flowing through that infrastructure [19].

**Specialized Test and Experimentation Infrastructures.** It is very difficult to experiment with Cyber C2 because of the scale needed to construct realistic testbeds and the need to isolate these testbeds in order to run malicious attacks. FIT has developed an infrastructure for Cyber C2 experimentation and testing called the Virtual Infrastructure for Network Emulation (VINE) [20], [21]. VINE simplifies the construction and management of large networks and isolates the networks from others, while maintaining a control backplane for data collection and experiment control.

### III. THE MIRA AGENT SYSTEM

MIRA is a software agent toolkit and runtime system for the development of distributed intelligent agent-based applications. From an architectural point of view, MIRA is a lightweight framework containing interfaces to a number of common services. Services are constructed to a well-defined interface specification, and a default implementation, or service provider, of each service provided with MIRA. Additional services and service providers can be programmed and used in addition to, or instead of, the default implementations so long as they satisfy the interface specification.

#### A. Design Goals

There are a number of goals that the MIRA agent system is designed to achieve.

**Agent System Service Modularity.** The goal of agent system service modularity is to insulate the agent execution environment, called the MIRA Environment, and its services from the agents. Agents are designed to interact with the service interface rather than interacting directly with the implementation of the service (see Figure 2). The purpose of the MIRA Environment is to provide agents with policy-regulated access to key infrastructure services. By constraining the interaction to the service interface, the agents in the system remain unaware of how that service is implemented and, therefore, can make use of any service implementation that fulfills the service interface. For example, the MIRA Environment has by default two messaging service implementations, one UDP and the other TCP. Agents sending messages in a MIRA environment have no knowledge of which implementation is being used.

There are two benefits to this approach. First, there is a software engineering benefit to developing service-provider implementations to meet a prescribed service interface. This simplifies implementations and ensures that service-provider implementations that are specific to a particular network installation or that make use of specialized equipment will be interchangeable with standard implementations. Second, there is a security benefit in not having all implementations for services in a network being the same. The idea of increasing the diversity of the network while maintaining functionality is central to the concept of moving-target defenses [22].

**Security limitations on agents.** MIRA agents have limited visibility into the lifecycle and service management functionality maintained by the MIRA Environment, so that agents cannot intentionally or unintentionally access agent system functionality that could be used to disrupt agent services. For example, agents have access to the agent discovery and lookup functionality of the Registration service, but they do not have access to register and deregister either themselves or other agents. By maintaining this separation between infrastructure functions and agent functions, we reduce the likelihood of deliberate or accidental errors caused by agents. Additional security measures include code isolation through the use of different class loaders for agents and infrastructure components, encryption, authentication, and constraints on code mobility. MIRA agents use a different class loader from the MIRA Environment class loader, so that malicious agents cannot load or access service classes by reflection. All infrastructure communications are encrypted by default, and

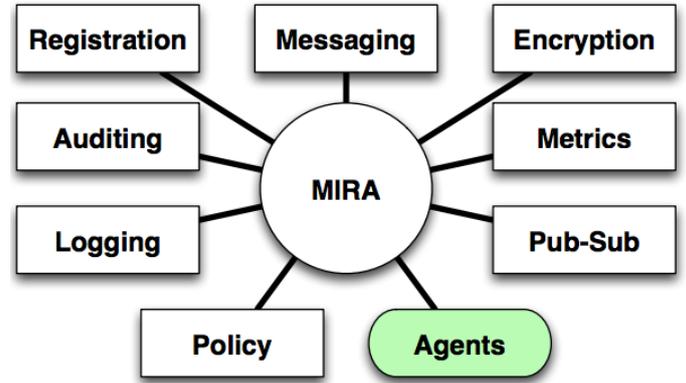


Fig. 2. On a given instance of the MIRA Environment, individual Agents communicate with Environment to access Core Services.

PKI-based authentication is used to ensure that all participants in communications can be identified. A decision has been made to explicitly prevent agent code mobility, so that all MIRA Environments must contain all of the agent code necessary to run any agent that moves into the Environment.

**Operating with Default Services.** When a MIRA Environment starts up, it consults a local configuration to determine which services and agents should be started immediately. If this configuration file does not specify a specific implementation for a service, a default service is used. The default service generally has a minimal set of configuration dependencies, and they are typically entirely self-contained. As is the desired behavior with all services, the default services either locate or create the service at startup, and negotiate with other existing components as necessary. Default services allow agent developers to focus their development efforts on the agent itself, with very little setup and configuration of the MIRA Environment required.

#### B. Core Services

The services that are currently specified for MIRA are shown in Figure 2 and include:

**Registration Service.** This service provides a naming and location service for MIRA environments and agents in the network. Upon startup, both environments and agents register their addresses and properties with the registration service. This information is updated when agents or environments are modified, and removed when they are terminated. The registration service also provides methods for environments and agents to lookup the location of other environments and agents based on their registration properties. There are currently three registration service implementations provided with MIRA, two that require external servers, and one that is self-contained.

**Message Service.** The message service is used by MIRA to send messages between agents and environments. Different implementations of the message service typically instantiate different transport mechanisms, such as the TCP- and UDP-based implementations that are currently provided.

**Message Encryption Service.** The message encryption service provides the means for encrypting and decrypting data and messages. By having the encryption service as a separate

service from the messaging service, any message encryption can be applied to any message transport service transparently to the user. There is one public-key encryption service provided with MIRA.

**Audit Service.** The auditing service captures and publishes a timeline of all of the events issued by the MIRA infrastructure for use either internally for supervisory reasons or externally for debugging and visualization capabilities. There are three auditing services provided with MIRA, two supporting output of events to external databases, and one self-contained.

**Logging Service.** The logging service is a utility service for developers and operators that specifies naming and location information for logs that may be generated on many different network hosts. There is one logging service provided with MIRA that enables configurable output of MIRA logs to the console, disk, or to a syslog service.

**Publish and Subscribe Service.** The publish and subscribe service provides the means for specifying the types of output agents and environments create, and the types of input that they need for processing. There is one simple peer-to-peer publish and subscribe service provided with MIRA.

**Policy Service.** The policy service provides the means for controlling the operation of the MIRA services on an interactive, detailed level. Policy control of the services ensures that each action taken by the MIRA infrastructure can be checked in the context of the current operating conditions, as specified in a user-defined policy. There are two policy services provided with MIRA, one semantic-based and the other based on XACML [23].

**Metrics.** The metric service provides the methods and functions for monitoring and reporting performance and resource usage measurements in the MIRA environment and for specific agent actions.

#### IV. AN ILLUSTRATIVE SCENARIO

As an example of using MIRA for Cyber C2, we developed a scenario involving an enterprise network containing subnetworks with access control lists (ACL) for several different departments, including Human Resources, IT Security, Research, etc (see Figure 3). Each of these subnetworks runs a behavior system that emulates the behavior of typical operators in each of the subnetworks producing background traffic for the experiment. An attacker machine was created behind its own router on a different network. The scenario runs under VINE in a private OpenStack cloud environment.

The VINE experimentation controller was used to deploy services, reset the scenario after each experiment, and to collect data results. The experiments were started and stopped using an Ansible [24] script on the experimentation controller. The script started the FTP traffic generators on the client nodes and the attacker scripts on the attacker node.

A C2Agent was responsible for sensing the state of the FTP services and controlling the moving target defense. To accomplish this, the enterprises used the MIRA agent system as a control interface for all sensing, messaging and coordination between enterprises. Both moving target defense and

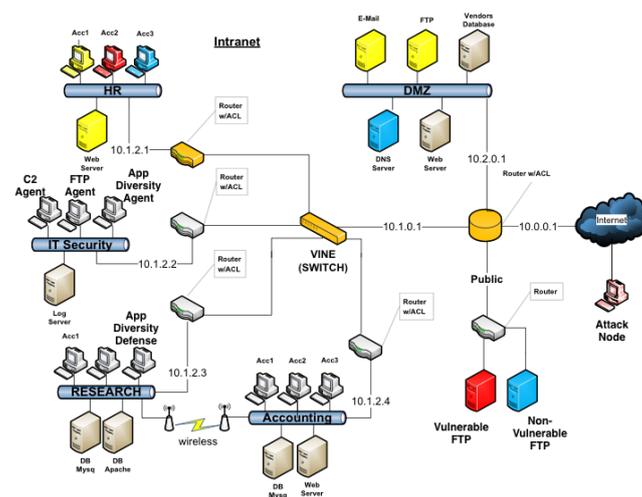


Fig. 3. Network diagram of the enterprise network used in the experiment.

service sensing was started automatically by MIRA, which launches a java virtual machine. This allowed the agent system to be started when the virtual machine started, which in turn started the moving target defense as well monitoring of the FTP service.

The attacker machine used a known vulnerability in one FTP application (proFTP) to terminate the FTP process and sleep for a predetermined time before performing the attack again. This behavior is meant to simulate a persistent attack from a malicious actor. The FTP servers in turn, had a watchdog task running that would restart the service when it was detected that the service was no longer running. The result of this attack and response is an FTP server that is intermittent and unreliable. The C2Agent will pick up on this diminishing service performance and will adjust the moving target defense configuration to favor the non-vulnerable FTP server in an effort to mitigate the persistent attack.

##### A. Agent Configuration

In the experiment, MIRA environments were running on several computers in the IT Security department (middle left of Figure 3.) The key agents were:

**C2Agent.** This agent coordinates and controls the behavior of the defense-specific agents by tracking overall system state and suggesting parameter and configuration changes for defenses to agents that interface to specific defenses. The C2Agent uses the Publish-Subscribe core service to subscribe to updates from application and defense monitoring and control agents that indicate the performance and status of the applications and defenses. It publishes control messages to application and defense agents to reset or modify the control parameters associated with the defense or application.

**C2 Interface.** This agent is a simple interface to the agent system that allows user interaction with the C2 Agent and defense-specific agents. Users can issue commands to the C2Agent or to the agents controlling the defenses directly. This agent runs in the same MIRA Environment as the C2Agent.

**AppOSDiversityC2Agent.** This agent is the proxy to the agent-based C2 system for the AppOSDiversity defense. The AppOSDiversity defense is a general purpose Moving Target defense that switches between several different applications that provide the same service. In our scenario, the defense switched between two applications that provide different instances of the FTP service, ProFTP and vsFTP (Very Secure FTPD). The AppOSDiversityC2Agent exposes to the C2 a number of controls of the defense, including the list of machines currently involved in the defense, and switching frequency between machines.

**FTP Agent.** This agent monitors the status of the FTP servers and reports the status of the FTP service to the C2Agent. It does this by attempting to log into the FTP service and retrieve a small file. The agent counts and reports how many times the retrieval is successful and how many times it is unsuccessful during a specified time window. The C2Agent can control this time window to enable faster, but noisier, updates or slower, and more accurate, updates.

Figure 4 shows the flow of messages between the agents used in this experiment. The principal function of the Cyber C2 is to perceive the status of the environment (FTP Agent sending FTP status to C2 Agent), present that information to the operator (C2 Agent sending C2 Status to C2 Interface), and then to make a decision and enact it (C2 Agent sending AppOS Command to AppOS Diversity Agent).

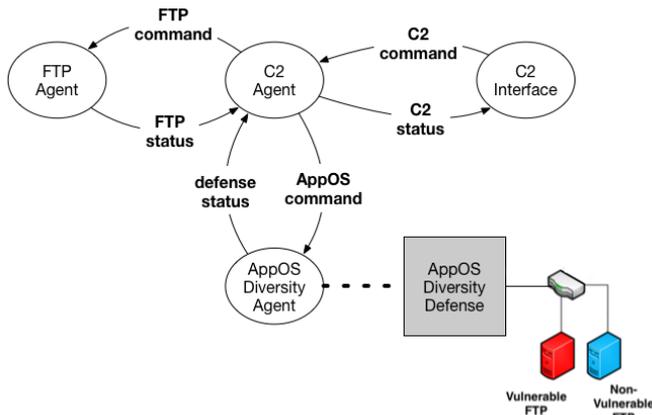


Fig. 4. Communication pattern of agents used in the experiment.

### B. Experiment Description

There were four separate experiments conducted and each was run multiple times to measure reliability and repeatability of the data collected. For all cases, only the subnet labeled IT Security in the scenario shown in Figure Figure 3 was used. They were:

**No Attack (Control).** While the defense was running, the attacker machine was turned off and the two FTP clients in each enterprise produced FTP traffic at a constant rate. The behavior system logged in, uploaded and then deleted a file from the server and measured the successes, failures, iterations and transfer metadata.

**Attack.** The attacker machine was activated and launched attacks on the public FTP site at regular intervals. Because a

	Dwell Percent	Attempts	Failures	Successful Requests
No Attacks	50%	306	4	98.69%
Attacks	50%	299	94	68.56%
Control Change	20%	311	76	75.56%
Further Change	10%	319	66	79.31%

TABLE I. SUMMARY OF THE EXPERIMENT RESULTS

moving target defense was in use, there was only a single FTP url and the attacking machine was unaware of which FTP it was attacking. Thus, the expected result would be that only half of the attacks and half of the FTP attempts would be successful. This is the baseline used as a reference to measure the improvements of a specific defense to the system.

**Control Change.** In this experiment, the C2Agent had the ability to dynamically reconfigure the moving target defense. An agent monitored the FTP application and log files and would trigger when it detected the FTP service success rate start to diminish. It would change the moving target defense settings to favor the FTP that was unaffected by the malicious actor. By configuring the moving target defense to use a 4:1 (20%) ratio instead of a 1:1 (50%), it is expected that the number of successful attacks will go down while the number of successful FTP attempts will go up.

**Further Change.** The C2Agent reduced the time spent in the vulnerable FTP service to 10% of the available time, where we expect to see further performance gains.

The results of the experiment are shown in Table I and Figure 5. Table I shows the "dwell percent", which is amount of time spent in each FTP application, the number of FTP attempts, the number of failed attempts, and the overall percentage of successful FTP connections.

Starting the attacks has a significant, immediate effect on the FTP success rate. With the AppOS Diversity defense running and spending an equal amount of time using each FTP server, the success rate is the lowest seen. This is because 50% of the time, the Cyber C2 chooses a vulnerable application to provide the FTP service. The reason the success rate is not 50% exactly is that the time to recover from an attack is slightly longer than the time to complete an FTP session. Therefore, more successful FTP transfers than failures occur during the same time interval.

Changing the control so that only 20% of the time is spent in the vulnerable service improves performance (Control Change), and reducing it to 10% improves performance still (Further Change). Eventually, the C2 will eliminate the vulnerable FTP application from being used, as long as the other FTP application continues to successfully respond to FTP requests.

## V. CONCLUSIONS AND FUTURE WORK

In this paper, we have motivated the need for proactive, adaptive Cyber C2, and presented several requirements for building one. We discussed our distributed communication and computation infrastructure called MIRA, and showed an example scenario run in our experimentation infrastructure, showing the effects of controlling a moving target defense control parameter.

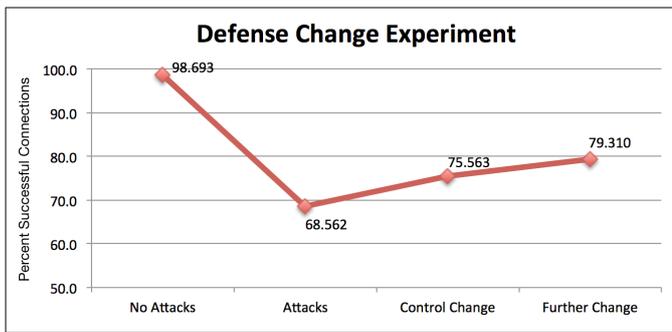


Fig. 5. Control changes to defense results in improved access to protected service.

Our next step for the research into resilient Cyber C2 systems is to quantify the effects of defense deployment, and the effects of changes in defense parameters on the mitigation of subsequent attacks. Additionally, we are expanding the scope of experimentation to include more complex and diverse attack models, as well as scenarios with incremental responses, such as terminating/redirecting/isolating connections, locking out/logging out users, killing processes, or quarantining messages.

Our next steps in the development of the MIRA infrastructure are creating additional providers, such as interfacing with Kafka as a Publish and Subscriber service provider, and using ZeroMQ as a Message service provider. We have also started the process of creating an open-source version of MIRA, OpenMIRA, that should be available soon.

#### ACKNOWLEDGMENT

This research project is sponsored by the U.S. Department of Defense. Any opinions, findings and conclusions or recommendations presented in this material are those of the author(s) and do not necessarily reflect the views of the Department of Defense.

#### REFERENCES

- [1] D. S. Alberts and R. E. Hayes, *Understanding Command and Control*. CCRP, 2006.
- [2] D. Andriess, C. Rossow, B. Stone-Gross, D. Plohmann, and H. Bos, "Highly Resilient Peer-to-Peer Botnets Are Here: An Analysis of Gameover Zeus," in *Proceedings of the 8th IEEE International Conference on Malicious and Unwanted Software (MALWARE'13)*. Fajardo, Puerto Rico, USA: IEEE Computer Society, October 2013.
- [3] M. Carvalho, T. C. Eskridge, L. Bunch, A. Dalton, R. Hoffman, J. M. Bradshaw, P. J. Feltovich, D. Kidwell, and T. Shanklin, "Mtc2: A command and control framework for moving target defense and cyber resilience," in *Resilient Control Systems (ISRCS), 2013 6th International Symposium on*, 2013, pp. 175–180.
- [4] M. M. Carvalho, J. M. Bradshaw, L. Bunch, T. C. Eskridge, P. J. Feltovich, R. R. Hoffman, and D. Kidwell, "Command and control requirements for moving-target defense." *IEEE Intelligent Systems*, vol. 27, no. 3, pp. 79–85, 2012. [Online]. Available: <http://dblp.uni-trier.de/db/journals/expert/expert27.html>
- [5] M. Bishop, M. Carvalho, R. Ford, and L. M. Mayron, "Resilience is more than availability," in *Proceedings of the 2011 Workshop on New Security Paradigms*, ser. NSPW '11. New York, NY, USA: ACM, 2011, pp. 95–104.

- [6] M. M. Carvalho, T. Lamkin, and C. Perez, "Organic resilience for tactical environments." in *BIONETICS*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, J. Suzuki and T. Nakano, Eds., vol. 87. Springer, 2010, pp. 22–29. [Online]. Available: <http://dblp.uni-trier.de/db/conf/bionetics/bionetics2010.html>
- [7] I. Kotenko, "Multi-agent modelling and simulation of cyber-attacks and cyber-defense for homeland security," in *Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications, 2007. IDAACS 2007. 4th IEEE Workshop on*, Sept 2007, pp. 614–619.
- [8] W. O. W. Group, "Owl 2 web ontology language document overview," W3C Recommendation, Tech. Rep., 10 2009. [Online]. Available: <http://www.w3.org/TR/owl2-overview/>
- [9] M. Atighetchi, N. Soule, R. Watro, and J. Loyall, "The concept of attack surface reasoning," in *The Third International Conference on Intelligent Systems and Applications (Intelli 2014)*, 2014, pp. 39–42.
- [10] M. L. Cummings, "Designing decision support systems for revolutionary command and control domains," Ph.D. dissertation, University of Virginia, 2004.
- [11] M. Carvalho, "A distributed reinforcement learning approach to mission survivability in tactical manets." in *CSIIRW*, F. T. Sheldon, G. Peterson, A. W. Krings, R. K. Abercrombie, and A. Mili, Eds. ACM, 2009, p. 21. [Online]. Available: <http://dblp.uni-trier.de/db/conf/csiirw/csiirw2009.html>
- [12] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, Massachusetts: MIT Press, 1998.
- [13] A. Servin and D. Kudenko, "Multi-agent reinforcement learning for intrusion detection," in *Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning*, ser. Lecture Notes in Computer Science, K. Tuyls, A. Nowe, Z. Guessoum, and D. Kudenko, Eds. Springer Berlin Heidelberg, 2008, vol. 4865, pp. 211–223.
- [14] P. Xie, J. H. Li, X. Ou, P. Liu, and R. Levy, "Using bayesian networks for cyber security analysis," in *Dependable Systems and Networks (DSN), 2010 IEEE/IFIP International Conference on*, June 2010, pp. 211–220.
- [15] M. M. Carvalho and C. Perez, "An evolutionary multi-agent approach to anomaly detection and cyber defense." in *CSIIRW*, F. T. Sheldon, R. K. Abercrombie, and A. W. Krings, Eds. ACM, 2011, p. 28. [Online]. Available: <http://dblp.uni-trier.de/db/conf/csiirw/csiirw2011.html>
- [16] L. B. A. Rabai, M. Jouini, A. B. Aissa, and A. Mili, "A cybersecurity model in cloud computing environments," *Journal of King Saud University - Computer and Information Sciences*, vol. 25, no. 1, pp. 63 – 75, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S131915781200033X>
- [17] T. C. Eskridge, D. Still, and R. R. Hoffman, "Principles for human-centered interaction design, part 1: Performative systems." *IEEE Intelligent Systems*, vol. 29, no. 4, pp. 88–94, 2014. [Online]. Available: <http://dblp.uni-trier.de/db/journals/expert/expert29.html>
- [18] P. Polack, M. Carvalho, and T. Eskridge, "Visualizing multi-agent systems," in *2013 IEEE/WIC/ACM International Conference on Web Intelligence*, 2013, Conference Paper.
- [19] T. Eskridge, M. Carvalho, P. Polack, F. Nebhard, and H. Thotempudi, "Interactive visualization of netflow traffic," Harris Institute for Assured Information, Tech. Rep. HIAI-TR-15-5-2, 2015.
- [20] E. L. Stoner, M. M. Carvalho, and T. C. Eskridge, "Vine: A cyber emulation environment for advanced experimentation and training," in *2014 National Security Innovation Competition Proceedings*, N. H. D. Foundation, Ed., 2014, pp. 14–17.
- [21] E. L. Stoner, "A foundation for cyber experimentation," Master's thesis, Florida Institute of Technology, 2015.
- [22] M. Carvalho and R. Ford, "Moving-target defenses for computer networks." *IEEE Security & Privacy*, vol. 12, no. 2, pp. 73–76, 2014. [Online]. Available: <http://dblp.uni-trier.de/db/journals/ieeesp/ieeesp12.html>
- [23] S. Godik and T. Moses, Eds., *eXtensible Access Control Markup Language (XACML) Version 1.0.*, February 2003.
- [24] (2015) Ansible home page, <http://news.ansible.uk/ansible.html>, retrieved on July 5th, 2015. [Online]. Available: <http://www.ansible.com/home>