

CyberCraft: Protecting Air Force Electronic Systems with Lightweight Agents

Lt Daniel Karrels
Air Force Institute of Technology
2950 Hobson Way
Wright-Patterson AFB, OH 45433
Email: daniel.karrels@aft.edu

Dr. Gilbert Peterson
Air Force Institute of Technology
2950 Hobson Way
Wright-Patterson AFB, OH 45433
Email: gilbert.peterson@aft.edu

Abstract—The United States military is seeking new and innovative methods for securing and maintaining its computing and network resources locally and world-wide. This document presents a work-in-progress research thrust toward building a system capable of meeting many of the US military's network security and sustainment requirements. The system is based on a Distributed Multi-Agent System (DMAS), that is secure, small, and scalable to the large networks found in the military. It relies on a staged agent architecture capable of dynamic configuration to support changing mission environments. These agents are combined into Hierarchical Peer-to-Peer (HP2P) networks to provide scalable solutions. They employ Public Key Infrastructure (PKI) communications (with digital signatures), and support trust chain management concepts. This document, a work-in-progress, presents the motivation and current challenges in choosing a network communications architecture capable of supporting one million or more agents in a DMAS.

I. INTRODUCTION

With the continued participation in the Global War on Terrorism (GWOt) by all sides involved, the level of technology employed by terrorists continues to increase, and threats to our electronic infrastructure compound daily. As a means of combating these threats, the United States Department of Defense (DoD) has been tasked to defend the United States Defense Industrial Base (DIB), to include all branches of military as well as resources essential to the functioning of those entities.

The scope of the DIB includes all elements necessary or utilized by organizations within the DoD. This includes, but is not limited to, portions of the internet (on top of which military medium and high encryption systems exist), DoD intranets and any external systems that interact with whom they interact, Supervisory Control and Data Acquisition (SCADA) systems in power, water, and waste water, wireless and radio systems used by military bases and military aircraft, and a host of other subsystems necessary to maintaining the functional defense of the United States, both at home and in deployed locations worldwide. These systems combine to form the electronics and communications infrastructure used by the DoD, are therefore subject to attack, and must be defended appropriately. The United States Air Force (USAF) has commissioned the creation of a tenth major command, the CyberWarfare Command, to address the defense and sustainment of these systems.

One of the focused research areas involved in the overall effort of the new USAF CyberWarfare Command is the development of a distributed multi-agent system, composed of lightweight software and hardware agents, to secure and sustain military networks and attached (including wireless) electronic systems. This project is called CyberCraft, and must:

- monitor systems and respond in near real time,
- respond according to current policies,
- provide feedback to human operators of alarms beyond a given threshold,
- support varying levels of autonomy, depending on situation and commander's intent,
- and be dynamically configurable during runtime.

In addition to these requirements, the system must also satisfy the austere restrictions for residing on a secured military network, both stateside as well as in more constrained deployed locations. The targeted number of agents is anticipated to exceed one million, encompassing workstations, servers, network routing devices, wireless systems, and other dedicated hardware systems.

The CyberCraft project is undergoing significant research in the military and defense contractor industries. It is centered around a lightweight agent that receives and executes arbitrary payloads as necessary, as well as monitoring and decision modules that are loaded dynamically and execute persistently in agent process space. The internal framework of the agent follows a multi-staged information flow model, with each stage capable of storing multiple (limited by memory constraints) dynamic modules. At present, these modules are organized into a behavior based system with limited local perceptual state. This permits both the activation and deactivation of behaviors dynamically as required by current policies, autonomous behavior, as well as providing a framework suitable for machine learning.

Standard modules also include communications capabilities for coordinating with other agents, as well as reporting relevant data to human operators. This paper presents an introduction to the problem domain, and several communications architectures that may be employed to support the project requirements. The architectures are primarily Peer-to-Peer (P2P) based, but

we are not limiting ourselves to this construct. The work in progress includes the design and execution of the experiments to quantify the suitability of each architecture to meeting project requirements.

II. MILITARY COMMAND AND CONTROL

The CyberCraft distributed multi-agent system (DMAS) will be used to defend existing and future military Command and Control (C2) systems. These networks carry with them additional requirements that separate them from business networks in the private sector. Namely, they:

- operate under tighter security constraints,
- require fault tolerance and self healing,
- and may affect human life.

In particular, the last attribute of C2 systems distinguishes them from most other systems. The capability to affect human life is a critical one to the mission of the military, and carries an extraordinary responsibility. Although CyberCraft will not have the ability to directly affect human life, it may do so inadvertently through inefficient communications, consuming too much network or processor load, or creating other scenarios that may deny real-time communications of critical data paths of C2 systems.

These requirements serve to create a separation between normal business class networks on which a typical DMAS system may operate, and the military networks on which the CyberCraft will operate. As such, the considerations found in much of the existing literature on P2P (and variant) systems dealing with distributed data storage [1] and federated search [2], [3] does not entirely meet the requirements of the CyberCraft.

III. PEER TO PEER NETWORKS

A P2P [4] network is one in which the nodes may establish multiple connections to other nodes. That is, the nodes are both client and server (or neither), and are free of the usual distinctions between the two. Rather, they communicate in a manner that best benefits the system objectives, without regard for the communication flow semantics of the client/server paradigm.

An extension to the P2P structure is the inclusion of super peers (or super nodes) [5]. These super peers act as regional hubs, absorbing additional burdens of the network traffic and processing load for distributed search and communications. These hubs may be interspersed across the network, as in the case of hubs controlling local clusters of regular nodes, or they may represent the bridges between layers of distinct P2P networks [6], [7]. Such networks are referred to as HP2P networks [8], [9].

HP2P networks provide additional scalability by linking multiple P2P networks together. The hubs that link the separate smaller networks may provide additional services (such as discovery), perform network management, secure communications by blocking certain types of messaging from passing between layers, or simply present a convenient point in the network to insert a human operator node for analyzing and

controlling the network. Although we use the HP2P architecture for scalability and compartmentalized security, many of the algorithms developed for strict P2P systems apply (some with modification) equally well to HP2P systems. Researching P2P systems is therefore a useful endeavor in creating large scale, secure, and reliable systems for the CyberCraft project.

IV. AIR OPERATIONS CENTER

The Air Operations Center (AOC) is the centralized control structure for waging an air war. Each theater of air battle is run by an AOC, and one exists and is responsible for each region of the world. The primary function of the AOC is to generate the Air Tasking Order (ATO) on a daily basis. The ATO is an electronic document describing the planned missions for a 24-hour cycle, including all armaments, aircraft equipment, targets, refueling, etc, necessary to achieve mission objectives. The AOC is perhaps the most important C2 structure the USAF possesses.

Several hundred to a thousand or more people work inside of an operational AOC, using upwards of 100 different pieces of software developed specifically to accomplish its mission planning and execution roles. Both security and sustainment of the network and software systems are critical to the success of the air mission. Providing these services to AOCs is one of the primary goals of the CyberCraft project.

The AOC is split into different groups based on roles. There are a dozen or more such groups, each tasked with a specific part of the AOC's overall mission. Each group uses different pieces of software and collaborates to operate on and generate data for injecting into the mission information flow. Each group implicitly describes a P2P structure, with the connections between groups forming a HP2P topology. The AOC is then connected to other operational systems of similar structure, further reinforcing the HP2P topology.

V. AGENT ARCHITECTURE

At its core, a CyberCraft is a software (and potentially hardware) agent capable of receiving payloads to deploy to the host machine. The agent architecture uses a staged module approach, where two or more module types are arranged linearly to support information flow (Figure 1). For our application, we currently support four stages, but the number of stages is not limited to any particular number. The implicit idea for this application is that the system is fundamentally a sensor network, with the added advantages of learning and large scale communications.

Stages in this architecture communicate by moving information from stage to stage, with modules receiving and processing this information as appropriate to their function and to the mission as a whole. Centered around the flow of information, this architecture implicitly supports flexibility and scalability.

For this application, the first stage supports any modules to initiate action or provide information to modules in other stages. The first stage of this architecture is envisioned as a sensor stage, collecting data about the agent's environment.




Fig. 1. CyberCraft Agent Architecture

Other modules for the first stage include communications modules. The second stage manages local perceptual state modules. As the application demands a small agent, the state tracking is minimal, based on the mission and policy the agent is executing.

Learning and decision making is performed in the third stage. For this application, the third stage is occupied by modules implementing the Unified Behavior Framework (UBF) [10]. This behavior framework is critical in that it supports simple and aggregate behaviors, and is designed to be modified at runtime. The UBF thus satisfies one of the more difficult requirements for this application: the ability to support commander's intent, which may change frequently and unexpectedly based on the needs of the mission and the agent's autonomy level.

The final stage used in the CyberCraft application is an actuator stage. For a purely software agent, this usually consists of communicating alerts and status information back to human operators. However, if the system is under attack or detects threats to mission or resources, such as a DoS or DDoS, modules in this stage may restructure the network or enable additional security constraints across the network.

VI. NETWORK OF AGENTS

When deployed to a network, these agents must organize themselves into a communications structure, and coordinate and cooperate to achieve mission objectives. The mechanisms for achieving the necessary information flow with minimal overhead are still under development. Message security is accomplished using the PKI, with a distributed key management system to protect the keys [11], based on Shamir key fragmentation [12]. The details of this protocol are in development and testing.

The deployment to the network will depend upon mission requirements. Since each agent autonomously functions on its own, the agents will be distributed based upon mission requirements. This is due to the fact that each computing platform may be integral to several ongoing missions in an organization. Given the simple autonomy afforded each agent, it is best if each agent only focuses on a specific mission, coordinating with the other agents on the same platform. In addition, the super-peer clusters of a hierarchical peer-to-peer network can be assigned based upon the mission. Where agents serving different missions may have different sensors,

behaviors, and actuators, agents within a certain role tend to have similar payloads. These agents communicate sensor and behavior information amongst themselves to provide time and distance based detection capabilities that are not possible with a single agent's information (such as Distributed Denial of Server (DDoS) and other attacks spread across the network or over a large period of time).

VII. EXPERIMENTAL OBJECTIVES

Toward the goal of scalability, the system is examined under several leading network communication architectures: hierarchical [13], P2P, and HP2P. In each case, a network of 1000 nodes is created in a testbed environment on the Air Force Institute of Technology (AFIT) clusters. At present, no more than 1000 nodes can be tested without them incurring penalties upon each other due to processing constraints. However, this number presents a significant step forward in efforts to improve the size of such networks.

We seek to examine the workload on each node and determine key bottlenecks in the network created by loading the network of agents to a point representative of an operational military squadron's network activities. At present, we are evaluating appropriate metrics to measure the suitability of each architecture to meet the requirements of defending and sustaining military C2 systems. With the exception of federated search metrics, these tend to fall into one of three categories.

A. Load Metrics

One possible metric measures the load on the network and the processing elements on which the CyberCraft functions [14]. Beyond a certain network utilization level, network latency increases exponentially with increased utilization [15]. This implies a necessary constraint on the communications protocols and frequency of those communications below a threshold to minimize performance impact. Such a constraint can be discovered by analyzing the traffic model of the target network(s).

Combined with the network load is the processing load. An agent's impact on its host machine can have a significant impact on the accomplishment of the host (and network's) mission. Enforcing reasonable maximum quotas on the processing power consumed by nodes in a DMAS is prudent and necessary to ensuring minimal mission impact. This threshold can be determined through trend analysis of the machines in each mission role, and the CyberCraft can be developed with this constraint in mind.

B. Cost Functions

This concept is divided into two approaches: cost functions based on measurable quantities, and those based on perceived importance of attributes. Models based on measurable quantities typically include things such as byte counts and frequency of communications. These can be combined with a traffic model tailored to the target network to yield a predicted cost of communications.

Models which use qualitative attributes rely on the experience and knowledge of network designers to choose appropriate measures to incorporate into the cost function. Each attribute is given a weight, based on its perceived importance. These weights can be somewhat arbitrary and reflect corporate knowledge of existing and future network use. Weights can be used to apply penalties to negative attributes like network diameter and bisection bandwidth, or to reward positive attributes such as number of links between hubs, etc. Note that some attributes can be used in both cost and reward functions. For example, it may be desirable to achieve a network diameter within a certain range based on previous experience or knowledge, and apply penalties for exceeding the maximum bound on that range, or rewarding a diameter less than the lower bound. This approach provides a lot of flexibility in network design, but does not present a provable means for meeting performance or mission related objectives.

C. Power Law Functions

During 1997-1998, experiments were performed to observe the traffic patterns at various points of the internet [16]. During this time, the internet grew in size by 45%, however the observations remained consistent through that growth. Scientists discovered that the structure of the internet closely followed a power-law relationship among several metrics: the diameter of the graph, the outdegree of any node, and the average outdegree of the nodes of the graph. Displayed together on a log-log plot, these attributes formed linear relationships. This yielded the notion of network topologies in large systems, in particular the internet, as following a power-law relationship. Three laws were observed and documented.

Power-Law 1 (Rank Exponent) *The outdegree, d_v , of a node v , is proportional to the rank of the node, r_v , to the power of a constant, \mathcal{R} :*

$$d \propto r_v^{\mathcal{R}} \quad (1)$$

Power-Law 2 (Outdegree Exponent) *The frequency, f_d , of an outdegree, d , is proportional to the outdegree to the power of a constant, \mathcal{O} :*

$$f_d \propto d^{\mathcal{O}} \quad (2)$$

Power-Law 3 (Eigen Exponent) *The eigenvalues, λ_i , of a graph are proportional to the order, i , to the power of a constant, \mathcal{E} :*

$$\lambda_i \propto i^{\mathcal{E}} \quad (3)$$

The configuration of the internet as a structure that obeys these power laws was perhaps an emergent behavior, resulting from necessity in growth, cost of hardware, and the vision of people involved in the decision making process. Whatever the origins, these power laws have been shown to adequately support large networks. As such, creating a network topology for the CyberCraft that follows these power laws is a reasonable approach, though challenging to create [17].

VIII. DESIGN OF EXPERIMENTS

We have developed a test bed capable running several thousand real agents. A small deployment platform following a subset of the CyberCraft agent requirements has been developed to live in this test environment. With this, we hope to perform small scale testing and data capture of the performance and fault tolerance properties of the various network architectures under consideration. In order to scale to larger numbers, we also require simulation of larger networks of CyberCraft, based on the small scale results, and mathematical extrapolation to networks in excess of one million nodes.

Generating the network configurations that obey both the power law distributions and minimize arbitrary and quantified cost functions will be done using an existing Multi-Objective Evolutionary Algorithm (MOEA) [18], tailored to appropriate metrics. This algorithm uses the Non-Dominated Sorting Genetic Algorithm (NSGA-II) [19] to determine minimum cost of multiple objectives, namely the cost and adherence to the three power laws and cost functions. The fitness functions of this algorithm are still under development, as we establish proper metrics to achieve the CyberCraft communication architecture objectives.

A. Traffic Models

We have chosen two traffic models from military networks, and seek a third from the private sector. The military traffic models used reflect the intended environments for the CyberCraft: a base network, and an AOC. The traffic found on a base is similar to that found in a private business of equal size, typically 1,000-20,000 users, as most users at Continental US (CONUS) bases carry out the same functions as normal businesses, such as word processing, email, and web browsing. The second military traffic model is the AOC, whose data flows are recorded specifically for research and testing. The pitfall of this traffic model is the security clearance necessary to view it, and this unfortunately corresponds to any analysis of the model as well. Therefore, even though we will perform experiments with this data set, the analysis will not be included in the final document.

A typical traffic model found in the private sector will be used as a baseline for our experiments. There exist a multitude of such traffic models, and we are current evaluating which is the most appropriate for our application and is accepted throughout the research community.

IX. FUTURE WORK

We expect to perform these experiments over the next few months, and to document results thereafter. However, we are open to input on our methodologies, chosen architectures, and assumptions, in hope of generating the most robust results possible.

REFERENCES

- [1] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proceedings of the 2001 ACM SIGCOMM Conference*, 2001, pp. 149–160.

- [2] J. Lu and J. Callan, "Federated search of text-based digital libraries in hierarchical peer-to-peer networks," in *proceedings of the 27th European Conference on Information Retrieval*, 2004.
- [3] —, "Content-based peer-to-peer network overlay for full-text federated search," in *Proceedings of the Eighth Recherche d'Information Assistee par Ordinateur (RIAO) Conference*, 2006.
- [4] I. T. Foster and A. Iamnitchi, "On death, taxes, and the convergence of peer-to-peer and grid computing," in *IPTPS*, 2003, pp. 118–128.
- [5] B. Yang and H. Garcia-Molina, "Designing a super-peer network," *icde*, vol. 00, p. 49, 2003.
- [6] H. Zhang, W. B. Croft, B. Levine, and V. Lesser, "A multi-agent approach for peer-to-peer based information retrieval system," *aamas*, vol. 01, pp. 456–463, 2004.
- [7] H. Zhang and V. Lesser, "A dynamically formed hierarchical agent organization for a distributed content sharing system," in *IAT '04: Proceedings of the Intelligent Agent Technology, IEEE/WIC/ACM International Conference on (IAT'04)*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 169–175.
- [8] H. T. Kung and C.-H. Wu, "Hierarchical peer-to-peer networks," Institute of Information Science, Academia Sinica, Taiwan, Tech. Rep. IIS-TR-02-015, April 2001.
- [9] L. Garcés-Erice, E. W. Biersack, K. W. Ross, P. A. Felber, and G. Urvoy-Keller, "Hierarchical p2p systems," 2003.
- [10] B. G. Woolley and G. L. Peterson, "Genetic evolution of hierarchical behavior structures," in *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*. New York, NY, USA: ACM, 2007, pp. 1731–1738.
- [11] S. Capkun, L. Buttyan, and J.-P. Hubaux, "Self-organized public-key management for mobile ad hoc networks," *IEEE Transactions on Mobile Computing*, vol. 2, no. 1, pp. 52–64, 2003.
- [12] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [13] V. Pappas, D. Massey, A. Terzis, and L. Zhang, "A comparative study of the dns design with dht-based alternatives," in *INFOCOM*, 2006.
- [14] B. Cooper and H. Garcia-Molina, "Studying search networks with sil," Tech. Rep., 2003.
- [15] T. G. Robertazzi, *Computer networks and systems: queueing theory and performance evaluation*. New York, NY, USA: Springer-Verlag New York, Inc., 1990.
- [16] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *SIGCOMM*, 1999, pp. 251–262.
- [17] C. R. Palmer and J. G. Steffan, "Generating network topologies that obey power laws," in *Proceedings of GLOBECOM '2000*, November 2000.
- [18] C. A. Coello Coello and G. B. Lamont, Eds., *Application of Multi Objective Evolutionary Algorithms*. World Scientific Publishing, 2004.
- [19] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II," in *Parallel Problem Solving from Nature – PPSN VI*, M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, and H.-P. Schwefel, Eds. Berlin: Springer, 2000, pp. 849–858.