

Business Network Fellowship Statement of Purpose: Distributed decision agents for managability and security

John Mark Agosta

March 31, 2006

1 Introduction

This application to the *Santa Fe Institute's* Business Network Fellowship presents a model for distributed control that applies to current and planned work in computer network manageability and security. First I review the problems facing the industry in networks of large numbers of machines. Second, I describe methods from a current project on self-propagating worm detection. Finally I embed the method used in this project in a larger theory of distributed decision-making agents, and describe how this fits in with current work at Intel. In short, this is a proposal to apply methods with origins in Bayesian decision theory and modeling to autonomous collaboration among networked computers.

2 Industry challenges

Intelligent, self-controlling computer networks—terms used in the industry are “autonomic systems”, or “embedded IT”—are still in very early stages. The desire for such technology comes from projections of the operations costs for computer networks, and the hope that automating these operations is feasible. The portion of IT organization resources dedicated to operations is 50 percent of the IT budget for personnel for installation, maintenance, inventory, upgrades, diagnostics, repair, and recovery. As a comparison, depreciating hardware is now only about 25 percent of total IT cost [4].

Typical current network operations consists of remote monitoring of event streams by human operators at consoles. There is some software support to categorize and filter events, but response is still largely by a human in the loop, and frankly, in our experience, most event data is simply dropped on the floor. Intel's current model to extend remote management to individual's desktops—promoted as *Active Mangement Technology* (AMT)—extends just this model by placing monitor and control capabilities in hardware within each platform, and does not address issues of intelligence or autonomy, arguably complicating the management burden.

There is some promising published research in intelligent self diagnosis [7] and configuration [9], in the spirit of the work I intend to pursue with the Institute.

3 A distributed detection scheme

As a member of an Intel project I have co-developed a scheme for distributed detection and inference (DDI) of malicious activity in computer networks, specifically to detecting self-propagating worms[1]. We've used a novel agent framework for worm detection, building a belief-passing inference scheme based on epidemic messaging that attempts to outrun the worm's epidemic propagation. DDI is a method for collaboration among individual nodes. Nodes can play two possible roles, which are not necessarily exclusive, one as a individual detector, the other as an aggregate detector. The local detector considers only the evidence available at an individual machine, and reports this to the aggregate detector. The aggregate detector infers the network state from the evidence it receives and raises an alarm if it concludes the network is under attack, terminates, or continues by propagating its belief to its neighbors. The inference scheme, expressed as a Dynamic (i.e. multi-stage, temporal) Bayes Network (DBN) maps

each stage in the DBN sequence to one node in the computer network. With analysis and simulation we've shown that unacceptable false positive rates for individual detectors can be reduced to acceptable system levels while also detecting slow worm traffic buried in normal background. This is a consequence of the law of large numbers applied to the individual detector signals (where "large" in this case means about 20).

It is advantageous for the nodes participating in DDI to randomize their messaging, so that malicious agents on the network cannot interfere with DDI messages, or falsify evidence that the aggregate nodes receive. Similarly malicious agents will randomize their behavior to try to avoid detection. Thus the messaging aspect takes on the aspect of a competitive game, with significant communication costs [6].

3.1 The sensor net problem

DDI can be viewed as a special case of a sensor network. In determining the network's state, a query of the network must trade off the cost of communication and sensing with the accuracy of the result. Query by message passing with sequential sampling is central to the approach of Deshpande et al.[5] These are the assumptions: The network monitors values that form a regular function of spatial extent (e.g. variations in temperature, in their example). The sensors measure a discrete, sparse sample of this function, which is modeled by the network. Querying the network, typically requiring combination of measurements from several sensors is implemented as a distributed observation plan. A plan consists of a circuit of nodes to visit and measurements to retrieve at each node. The observation plan is "injected" into the network where it executes on individual nodes by exploiting their local information as it travels through the network. The plan may compute its strategy as it proceeds to trade-off costs of collecting the measurements with the value of information gained. The plan is analogous to an "information collection agent" that travels from node to node.

3.2 Roaming belief agents

Generalizing DDI, an agent may work like this: In a network of nodes—each with sensing and computational capabilities, and limited communication abilities—the "agent" consists of a combined computation and messaging scheme. The agent is "born" at a node, perhaps to meet an external need, in the form of a belief about its surroundings, limited by the node's perceptions. This belief moves to an adjacent node that incorporates its belief with the belief it has received, which then propagates the updated belief to additional nodes. This propagation of belief scheme fits naturally into a sequential inference/decision making "MDF" style method, assigning each stage in the computation to one node in the sequence. At each stage, the agent can infer, choose a propagation step, or act, based on its belief.




Figure 1: An agent as a distributed inference computation. The agents in this model are identified with a distributed, sequential decision calculation, e.g. a Dynamic Bayes Network that moves from node to node. Each node is capable of receiving a message, and updating its belief about the world state contained in the message with its local information. The node then takes an action, depending upon its belief; one of the actions being to transmit its belief to a neighbor who continues the sequence. This computational sequence of messaging and inference constitutes the "agent."

The agent computation is described by an objective, a probabilistic model of the world, and a set of actions to maximize the objective based on the beliefs obtained in the world model. Consider a random walk in the belief's state space, driven by the information updates. The agent takes actions depending on the area occupied by the belief random walk. Thus two paths are being traced out simultaneously by the agent; its itinerary through the physical network and the trajectory of its belief in belief state space.

Naïve Bayes is an example of a local belief update model. We model the state space of beliefs from the agent’s perspective. With each observation the agent updates her belief, taking a step in belief space. Bayesian update provides the update method. The distribution over messages as generated by the agent’s environment results in a distribution over steps in belief space, which can be viewed as a *random walk*, a well studied field.[8] By assigning actions to boundaries in the space, the model of the agent can be completed, and her behavior analysed. The model can be generalized by making step size a function of time or previous state, capturing the dependencies when naïve Bayes is supplemented with the dependencies of a general Bayes network.

3.3 Agent decision-making

We consider a specific class of agents that implement a model of sequential decision-making. In summary, the characteristics of the agents:

- Agents consist of a physical probabilistic model, a preference model and a set of allowable actions. Actions include information collection and roaming (messaging) options.
- Agents have private knowledge, described by a distribution of probability over the states in the world that they track. This constitutes a belief state.
- They observe local information, at a cost, which together with their model of the world makes it possible for them to update their belief. Maintaining belief is interpreted as probabilistic inference, and adaptation can be viewed as inference about inference, a second order inference process. Probabilistic inference is interesting when observations, communication and effects of actions are unreliable and possibly costly.
- Belief is communicated, also at a cost, to neighboring nodes. Since belief summarizes relevant aspects of the world, it minimizes communication costs, compared to transmitting information data among neighbors. Agents take action based on their beliefs. From the node viewpoint, decision-making is distributed, again taking advantage of the parsimony of messaging with belief.
- Multiple agents in the network may all be on the same team by means of shared objectives, and possibly face a malicious adversary or just dis-interested natural hazards.
- Agents are not reactive but optimizing and intensional. We gauge performance by how well local pursuit of goals is able to achieve a global goal. This leads naturally to a concept of comparative efficiency, by comparing the distributed system’s performance to what can be achieved by central control. In unreliable systems, distributed performance may exceed centralized, as we have observed in DDI simulations.

4 Research Objectives

Research done under this fellowship would extend current efforts in centralized models[2, 3] to the distributed framework presented here. Specifically, approaches to diagnostic and predictive inference and its role in manageability would benefit from this work. With the industry’s interest in intelligent “self-healing” networks, and the advertised capabilities of upcoming platforms to engage in distributed functions, there are numerous uses to which such agent models may be applied. For instance, Intel’s announced AMT processing capability, independent of the main CPU, is to be incorporated in each chip. We face the opportunity of continuing growth of local machine processing power, but partitioned among multiple cores and virtual machines. The coördination of multiple processing units, not only for parallel computation, raises the more interesting prospect—in the context of this research—of more reliable and intelligent systems that would require solving some of the challenges this research addresses.

4.1 Methods and tools

The basic modeling techniques that this work draws upon can be traced back to early work in decision-theory, Markov Decision Processes and the variety of sequential methods that exploit Bellman's equation. Currently there are numerous implementations in the Machine Learning, Uncertainty in Artificial Intelligence and Graphical Models community of research tools. In our current project we've taken advantage of the communities' software tools, and gone as far as to duplicate some of them in open source libraries that Intel has developed such as PNL. (<https://sourceforge.net/projects/openpnl/>). The experimental testbeds from the DDI project are also a natural starting point for further investigation.

4.1.1 Author Bio

John Mark Agosta is part of Intel Research's DDI network intrusion detection project. His work involves developing probabilistic models for intelligent diagnosis, and management. Previously he worked on automated response to customer inquiries for Customer Relationship Management (CRM) software while working at Edify Cororation. From 1998-2000 he was Chief Technical Officer (CTO) for Knowledge Industries, where he built Bayes networks for medical, avionics and automobile clients. From 1992 to 1998 he worked as a research engineer at SRI International. He built models for electric utility generator alarm filtering and computer network intrusion detection. He also worked in automated planning for emergency response and USAF air campaign planning. Agosta received his Ph. D. in the Engineering-Economic Systems Department (now Management Science and Engineering) of Stanford University in 1991. His thesis topic was on an application of Bayes networks to visual recognition.

References

- [1] Agosta, J.M., D. Dash, E. Schooler, B. Kveton, Distributed Network Attack Detection, Adaptive and Resilient Computing Security, Santa Fe Institute 2-3 November (2005).
- [2] Agosta, J.M., T. Gardos, "Bayes Network 'Smart Diagnostics' *Intel Technology Journal*. http://developer.intel.com/technology/itj/2004/volume08issue04/art10_bayesnetwork/p01_abstract.htm (November 2004).
- [3] Agosta, J.M., J. S. Katz "The use of Evidence Conflict to extend Diagnostic Models" in Kai Goebel and Piero Bonissone, Cochairs, "Information Refinement and Revision for Decision Making: Modeling for Diagnostics, Prognostics, and Prediction," Papers from 2002 AAAI Spring Symposium (AAAI, Technical Report SS-02-03, 2002). <http://www.aaai.org/Press/Reports/Symposia/Spring/ss-02-03.html>
- [4] Busch, D., Bryant, G., Sayles, B., Swinford, T., "The Digital Office: Cross-Platform Embedded IT for Manageability, Security, and Connectivity," *Technology@Intel Magazine*, Sept. (2004).
- [5] Deshpande, A., C. Guestrin, S. R. Madden, J. M. Hellerstein, W. Hong, "Model-Driven Data Acquisition in Sensor Networks" *Proc. 30th VLDB Conference* (Toronto, CA, 2004).
- [6] Emery-Montemerlo, R. "Game-Theoretic Control for Robot Teams" Doctoral dissertation, (Tech. Report CMU-RI-TR-05-36, Robotics Institute, Carnegie Mellon University, August, (2005).
- [7] Littman, M. L., T. Nguyen, H. Hirsh, E. M. Fenson and R. Howard "Cost-Sensitive Fault Remediation for Autonomic Computing" (IJCAI 2003).
- [8] Spitzer, F. *Principles of Random Walk* (NJ: Van Nostrand, 1964).
- [9] Tesauro, G., D. M. Chess, W. E. Walsh, R. Das, A. Segal, I. Whalley, J. O. Kephart, S. R. White, "A Multi-Agent Systems Approach to Autonomic Computing" in *Third International Joint Conference on Autonomous Agents and Multiagent Systems - Vol. 1 (AAMAS'04)*, pp. 464-471, (2004).