

# Infrastructure web: Distributed monitoring and managing critical infrastructures

Guofei Jiang, George Cybenko And Dennis McGrath\*

Institute for Security Technology Studies  
Thayer School of Engineering, Dartmouth College  
Hanover, NH 03755, USA

## ABSTRACT

National-scale critical infrastructure protection depends on many processes: intelligence gathering, analysis, interdiction, detection, response and recovery, to name a few. These processes are typically carried out by different individuals, agencies and industry sectors. Many new threats to national infrastructure are arising from the complex couplings that exist between advanced information technologies (telecommunications and internet), physical components (utilities), human services (health, law enforcement, emergency management) and commerce (financial services, logistics). Those threats arise and evolve at a rate governed by human intelligence and innovation, on “internet time” so to speak. The processes for infrastructure protection must operate on the same time scale to be effective. To achieve this, a new approach to integrating, coordinating and managing infrastructure protection must be deployed. To this end, we have designed an underlying web-like architecture that will serve as a platform for the decentralized monitoring and management of national critical infrastructures.

**Keywords:** Infrastructure protection, cyber security, architecture, monitoring and management, distributed system

## 1. INTRODUCTION

Modern threats to critical national infrastructure are evolving at the same rate as the technology on which that infrastructure is based. This is a key axiom of the work described in this paper. To illustrate the point, consider the following chronology of events related to the recent Distributed Denial of Service (DDOS)<sup>1</sup> attacks launched against major e-commerce companies.

Early summer of 1999	DDOS capabilities are demonstrated at a European “hacker festival.”
Late summer of 1999	First DDOS attacks at the University of Minnesota are detected and documented.
November 1999	A workshop on DDOS attacks and defense mechanisms is hosted by the Computer Emergency Response Team (CERT) <sup>2</sup> , Carnegie Mellon University.
December 1999	Programs for detecting DDOS “zombies” are distributed.
February 2000	DDOS attacks are launched against major internet sites.
March 2000	The possibility of a DDOS-type attacks against the 911 system is identified.
April 2000	DDOS-type attacks against the 911 system are suspected in Texas.

---

\* Authors' email addresses: Jiang: [gjf@dartmouth.edu](mailto:gjf@dartmouth.edu); Cybenko: [gvc@dartmouth.edu](mailto:gvc@dartmouth.edu); McGrath: [dennis.mcgrath@dartmouth.edu](mailto:dennis.mcgrath@dartmouth.edu)

Sometime in the future DDOS attacks within the financial sector, using automatically generated consumer trading, are detected.

This chronology illustrates three major points:

- a. The time intervals between when a new threat is identified, when it manifests itself and when it is modified (mutated) into different forms are relatively short and appear to be shrinking;
- b. Cyber security and physical security are rapidly becoming intertwined, and attacks against information systems can have immediate and profound effects on the physical systems that depend on them.
- c. Threats within one sector (telecommunications/internet) can easily spill over into other sectors such as human services (the 911 system) and the financial system.

To meet these challenges, we need to leverage modern information technologies and create an infrastructure protection process that can operate seamlessly at an accelerated time scale. Moreover, that process must be able to monitor and manage the complex interactions between infrastructure segments that are becoming the norm. This is especially important considering the fact that many recent attacks on national infrastructure have been credited to pranksters and individuals working alone or in small groups. We have not yet really seen what kinds of damage well-financed, coordinated, professional attacks are capable of creating.

Like the World Wide Web, the Infrastructure Web should have the following characteristics:

1. It should be decentralized, asynchronous and redundant;
2. New elements can be added to it or old elements can be removed from it by authorized personnel but without centralized control;
3. It should be searchable and self-organizing;
4. It should allow new services to be built easily on top of existing services;
5. It should allow for multiple, redundant communication paths between entities.

Section 2 of this paper describes the various stages in the Critical Infrastructure Protection process today together with our vision for how those stages can be integrated. Special attention is given to infrastructure related to information technology, namely internet and telecommunications, but we indicate how the ideas can be generalized to other infrastructure segments.

Section 3 describes the conceptual organization of the Infrastructure Web that we are currently implementing. The functional operation is illustrated through some examples. Meanwhile section 3 also gives a brief technical description for how the Infrastructure Web can be implemented, using current computing and networking technologies.

Section 4 discusses some related work, and Section 5 is a summary.

## **2. THE INFRASTRUCTURE PROTECTION PROCESS**

The emergency management, public health and more recently computer security communities have decomposed their management processes into smaller, logically-concise stages. For example, the DARPA Information Assurance program is using the three-stage “Protect-Detect-React” paradigm to organize work within that area.<sup>3</sup> Figure 1 shows the six stages we propose for Information Infrastructure Protection. These stages roughly correspond to stages used in other emergency management areas with different degrees of granularity perhaps. We briefly describe each stage and its relationships with other stages.

# Information Infrastructure Protection

## Stages and Process

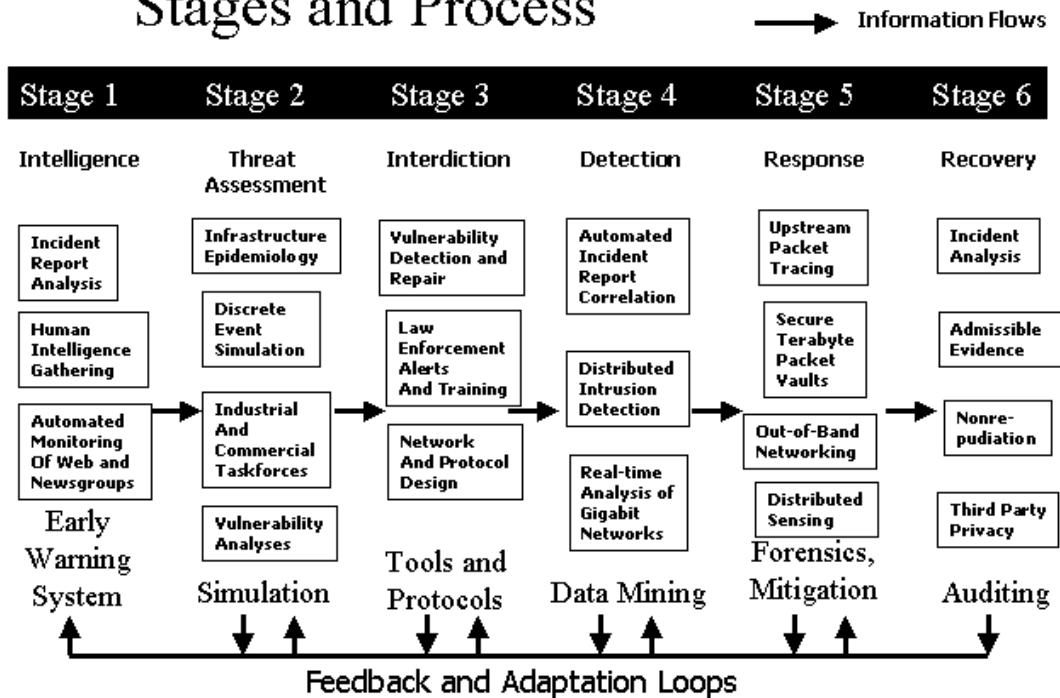


Figure 1: The Information Infrastructure Protection Process

### 2.1 Intelligence

The first step in infrastructure management is intelligence gathering about emerging threats. This is typically done using human intelligence reporting, analysis of unusual incidents, and information harvesting from open sources such as the web and news sources. This is the early warning system that can identify new threats early in the process, before they manifest themselves in real attacks or disasters. “Red teaming,” namely the use of selected experts for scenario building and threat design for proactive analysis, is an important part of this stage. We include that in the “human intelligence” component.

Figure 2 identifies three sources of intelligence for early threat identification: incident analysis, human intelligence and automated tools for harvesting and organizing information from open sources such as the web and newsgroups. In the information infrastructure protection problem, early evidence of threats are often proposed and discussed in such open sources. Such open sources are useful for human-initiated threats but not so useful perhaps for predicting complex interactions between infrastructure elements or natural events and design flaws. Human intelligence is more important for identifying those threats.

From the point of view of automating this stage of the process, automated incident report analysis and monitoring of open web- and internet-based sources are most promising. Several organizations already provide on-line access to incident reports and threat alerts (see <http://www.cert.org> for example) although those resources are not organized to allow powerful search capabilities through a database engine interface.

Ideally, a new incident report could be quickly and automatically matched against an on-line database of previously seen threats and attacks to see if the threat is novel or known. Today, this stage of early warning is done by experts who rely on their own memories, networks of colleagues and ad hoc searches of archives of previous attacks.

Automated monitoring of the web and various news groups for early threat identification is technically possible today<sup>4</sup> but not done to our knowledge. We are currently developing such a capability.

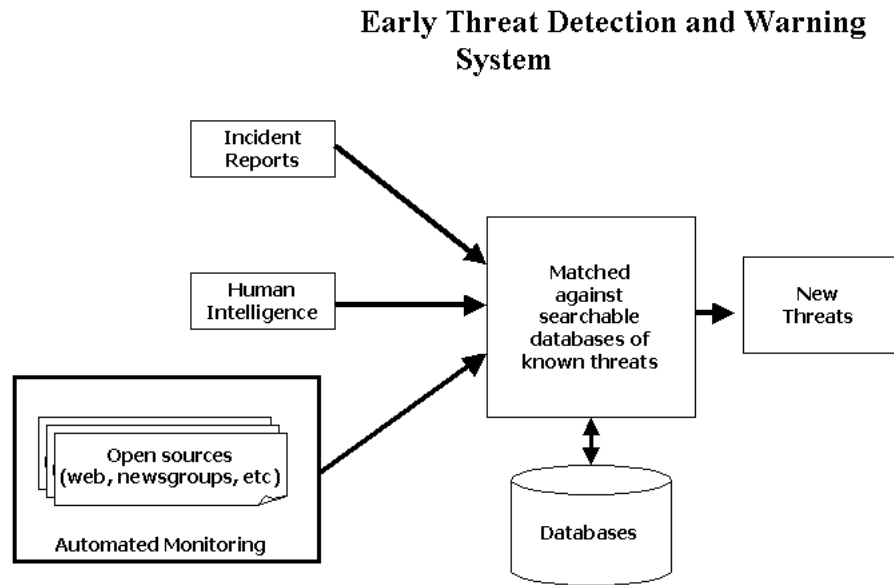


Figure 2: Early Threat Detection and Warning

## 2.2 Threat Assessment

Once a new threat is identified, risk assessment and some sort of “cost-benefit” analysis of responding to the threat must be performed. This stage requires some sort of epidemiological model of how an attack or failure based on the threat will manifest itself and how it will affect other infrastructure systems. Basically, the question is: what its dynamics are? Related to this of course are the costs associated with containment or interdiction versus the costs of an attack or failure based on that threat. As is often the case in defensive strategies, the cost of defense can be much higher than the cost of the attack, but that must be weighed against the social and human cost of major systems failures.

At present, our understanding of “infrastructure epidemiology” is very poor, at least in the open literature. The challenge here is to develop quantitative models of how vulnerabilities are distributed nationally or even globally, and how failures based on those vulnerabilities can cascade through the overall infrastructure.

Such analyses will probably have to rely on large-scale discrete event simulations since closed-form solutions are highly unlikely. Government, industrial and commercial task forces must be able to provide quick and reliable input into the vulnerability assessment process so that some form of realistic cost-benefit analysis can be performed in the threat assessment stage.

## 2.3 Interdiction

The interdiction stage of infrastructure protection attempts to proactively prevent or prohibit attacks or failures based on known threats. Virus scanners, software patches and improved network designs and protocols are examples of interdiction in the information infrastructure segment. An important element of interdiction is the training of system operators and law enforcement personnel, especially at the state and local levels since these communities are typically the first responders to attacks and failures.

These ingredients in the interdiction stage typically operate at different time scales. For example, the deployment of more robust and secure designs and protocols can take many years to permeate the infrastructure because of lock-in effects. On the other hand, software patches and virus scanning updates occur on the time scale of weeks quite often. The training of early responders such as system operators and law enforcement and emergency management personnel is problematic because of the huge demands on time and expertise in those sectors. The rate at which new threats and vulnerabilities are arising outstrips the ability of such personnel to attend training meetings and courses so that remotely accessible, distance training using networked interactive material is necessary. Cost-benefit analysis is essential to identify threats and vulnerabilities that are most likely to have high impact because existing time commitments and obligations preclude the ability of first responders to be prepared for all possible failures. This stage of the process must focus on interdiction in high-cost and/or high-probability events.

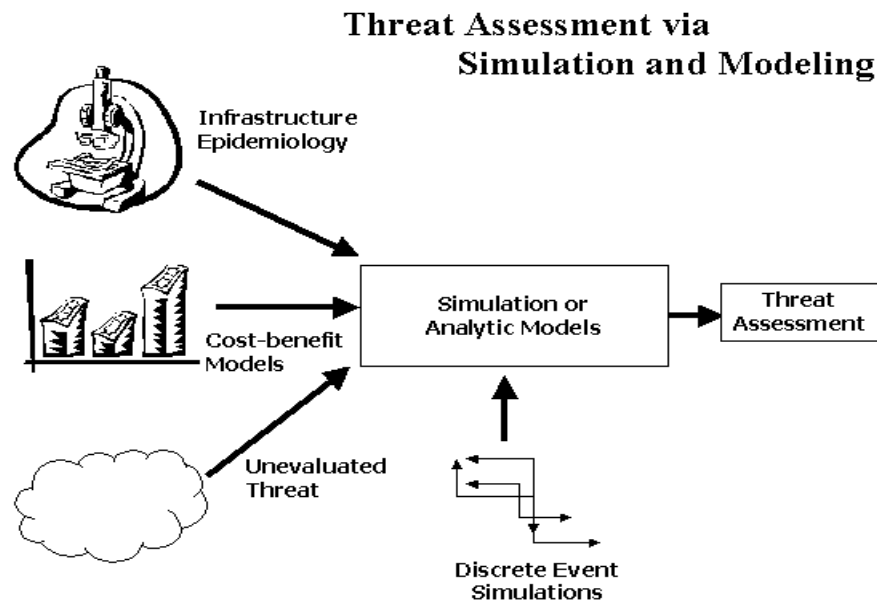


Figure 3: The Threat Assessment Stage

## 2.4 Detection

The detection of actual failures or attacks is enabled by monitoring distributed “sensors” that are positioned throughout the infrastructure itself. Raw sensor data must be harvested, mined, correlated and otherwise analyzed. Examples of such sensors include computer network monitors (based for example on SNMP agents or packet analyzers), public health records, medical laboratory results, environmental monitoring stations, financial market trend monitors and so on. Human observations in the form of natural language reports are also relevant to this stage.

Whereas the Early Warning System part of the process is meant to anticipate attacks and failures through proactive intelligence gathering and analysis, this stage is meant to respond to mature attacks and imminent failures. Ideally, threat assessment and interdiction has prepared the community for these events but may not always be the case.

The challenge in automating this stage of the process lies in flagging anomalous events that have not been seen before without generating large numbers of false positives. This requires training an automated system on “normal” and known behaviors, flagging behaviors that fall outside this regime. The technical challenge here is that many new behaviors emerge in the course of natural, non-threatening operating modes. Much work remains to be done in this area.

## 2.5 Response

Once an attack or failure has been detected, an appropriate response is required. We focus on law enforcement or internal auditing responses to information infrastructure events. A major challenge in responding to cybercrime and cyberterrorism attacks is identifying the source of the problem. This requires forensic techniques that allow building a trail of legal evidence for future investigation while respecting the privacy of third parties. These considerations require the ability to do fast and reliable upstream packet tracing, something that currently requires time consuming and relatively slow operator intervention. Moreover, the fact that many internet links are now operating in the multiple megabit and even gigabit per second range, archiving network traffic for forensic analysis is a major technical challenge. Early work in this area is promising but much development remains to be done.

Another fundamental challenge in responding to infrastructure failures and attacks is that the very systems, namely the telecommunications networks, that responders will have to use to coordinate a response are themselves part of the infrastructure and highly vulnerable to failures. Any future infrastructure web architecture must provide for “out-of-band” and otherwise redundant communication capability.

This can be accomplished through the use of multiple communication channels based on different protocols implemented by different vendors so that a single vulnerability does not compromise the whole system. In this case, standardization is bad for survival and we need heterogeneous systems. Additional out-of-band communication capability can be achieved by radio and satellite networking which is currently being investigated on several fronts.

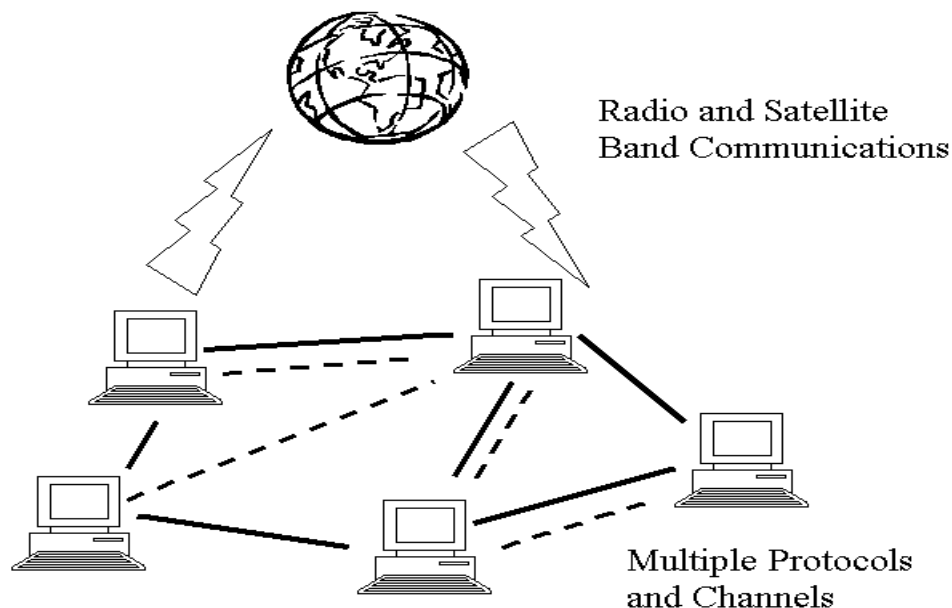


Figure 4: Out-of-band and redundant communications channels

## 2.6 Recovery

In the law enforcement arena, recovery from an attack or other criminal activity related to national infrastructure includes archiving non-reputable evidence without violating privacy laws and standards. Complete analysis of the incident is required to learn from it and to archive its characteristics in appropriate databases for future use in detection and training. Technical challenges here include training of first responders on the appropriate forensic techniques that accomplish these goals.

### 3. ARCHITECTURE

According to the report of the President's Commission on Critical Infrastructure Protection<sup>5</sup>, the infrastructure networks of greatest importance to national security and stability include telecommunications, electrical power systems, gas and oil, banking and finance, transportation, water supply systems, government services, and emergency services. To effectively protect these critical infrastructures, it will be necessary to have a system in place that can monitor and manage these very large, complex and dynamic networks. Our proposed Infrastructure Web system provides such a basic architecture, as well as the underlying paradigms by which a problem of this scale and scope can be addressed.

The above section has discussed the various stages in the critical infrastructure protection process and our vision for how those stages can be integrated. So now the question is: how to integrate and implement these stages and visions into a real monitoring and management system? We propose that national Infrastructure Web networks be built up with four types of basic distributed components: Directory service, Infrastructure server, Sensor web and Emergency Information Search server. All these distributed components will be organized and integrated throughout the national wide networks with Sun's Jini system.<sup>6</sup> Jini is designed for deploying and using services in a network and enables the construction of dynamic, flexible, and robust systems from independent distributed components. Further, access to Jini's source code has fostered a thriving community of developers who continue to enhance and expand Jini's capability. A framework of the infrastructure web architecture is shown in Figure 5. With this kind of architecture, we believe that the Infrastructure Web system can be exploited as a platform to implement our distributed infrastructure assurance vision.

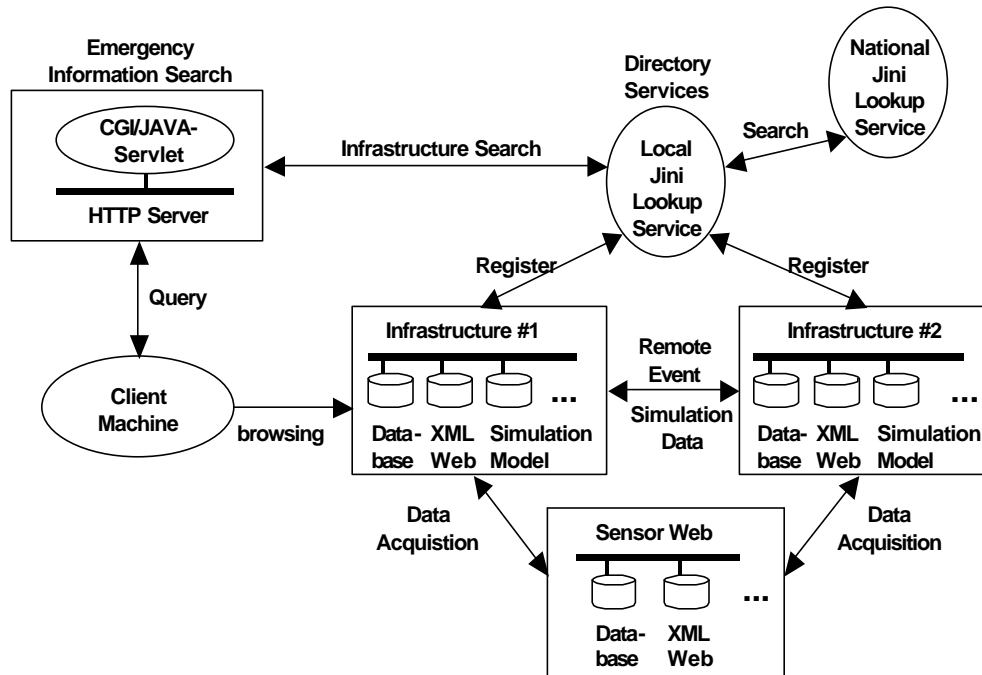


Figure 5: The architecture of the Infrastructure Web

Currently several efforts are underway in US to build the computational grids, such as Globus<sup>7</sup> and DARPA CoABS.<sup>8</sup> These projects are developing the fundamental technology that is needed to build the computational grids, execution environments that enable an application to integrate geographically distributed instruments, sensors, data products, displays and computational and information resources. Here so-called *grid computing* refers to computing in a distributed networked environment in which computing and data resources are located throughout the network. The vision is that these grid infrastructures will connect multiple regional and national computing power and data resources that support dramatically new classes of applications. In this sense, the infrastructure web system is actually a national wide grid that is specially developed for the critical infrastructure protection. We believe that with more and more sensors and other services on the infrastructure servers come to online, the infrastructure protection grid will grow quickly. Currently we are investigating the possible standard and transparent data exchange format, communication protocol, user interfaces and APIs (application programming

interface) among these infrastructures. So later by the integration with these existing distributed services, some new applications or services can be easily implemented for infrastructures' protection and then serve as new "existing services" for the further development of other applications or services. As the result, the scale of the grid is growing.

### 3.1 Infrastructure Server

In the Infrastructure Web system, one infrastructure network server represents one critical infrastructure element (such as a hospital or power plant) in the physical world, and the server's IP address is the unique identification for the infrastructure elements. Basically, the infrastructure server will have a real time database, an XML-based web, a simulation model and other infrastructure protection services running on various ports of a single server.

The database acquires real time data from the sensor web or other sources (e.g. host-based detection systems). These data consist of the security status, internal state information, and so on. Some data will be displayed on the web in real time to show the infrastructure's current status, and some will be used as input to simulations for automated evaluation (such as the threat assessment). Currently we employ a Mysql database server in our implementation because it is a free open source database server, and it is supported on multiple platforms. Open Database Connectivity (ODBC) and Java Database Connectivity (JDBC) are two of the most popular database interfaces that make database access transparent. Applications written in modern programming languages such as Java, Visual Basic and C++, can easily use the ODBC or JDBC function from the database driver to interact with remote databases. This kind of open database connectivity can be used as a standard data exchange approach between the distributed applications or services. On every server's web pages, the relational database server type and the table layout will be described, so a remote client can download the driver for that type and use ODBC or JDBC functions to call a SQL query to access the database. The advantage of this data exchange approach is that specific knowledge of the various communication protocols and programming interfaces is not needed for the client. Once the remote client knows the table layout for the database, he knows how to access the data. Moreover, the security policy of the database server uses multi-level security and authorizes remote clients different data access permissions.

It is widely accepted that extensible markup language (XML) is the most promising Internet technology development since Java. XML is a license-free, platform-independent language for describing and publishing structured data in text form. Where HTML only allows users to see formatted text, XML allows users to understand and use self-describing and structured data. For each of the eight infrastructure categories listed in the President's Commission on Critical Infrastructure Protection, an Infrastructure Markup Language (IML), created with XML document type definitions (DTDs), will clearly describe the attributes of the category in standard formats. XML is also an excellent vehicle for the interchange of data among different applications. By browsing an infrastructure server's XML pages, clients can check one or many state variables of the critical node, such as the packet traffic throughput of an important LAN infrastructure. The relationships between related infrastructures will be described by XML's X-Links and X-Pointers.<sup>9</sup> Meanwhile other web technologies such as Java applets and JavaScript will also be used to describe the infrastructures.

An essential part of the infrastructure server consists of appropriate analytical and simulation models of the infrastructures together with a description of their behavior under dynamically changing interconnections. Just as pins are wired in chips on a PCB (printed circuit board), large-scale discrete-event simulations can be implemented for threat assessments by "wiring" infrastructure models' input and output. We believe that with some adaptive learning technologies such as neurocomputing and evolutionary computing,<sup>10</sup> improved planning, control, and coordination strategies can be derived with simulations. These strategies and policies will help these related infrastructures cooperate efficiently once a disaster or attack occurs.

Other infrastructure protection applications and services will be implemented on the infrastructure server based on the specific requirements of the infrastructure element category. F.g. the early warning system can be implemented as an information infrastructure in our architecture to offer automated intelligence collection services. But the basic development process for these applications will be similar: search and collect the distributed sensor and infrastructure data that is related to that infrastructure; develop the application program to process these data and then to verify whether the infrastructure is running well or under threat (which needs specific knowledge support from experts in that area); if not, take steps in the application programs to respond and recover the possible failure. Meanwhile, if the application is going to offer data to other applications, it should also report the data with the standard user interface and API as we discussed.

### 3.2 Directory Service

The Infrastructure Web will have two directory service levels: a local/state level directory service and a national level directory service. All of these directory services will be implemented with Jini lookup services. Infrastructure elements need to register themselves in the local Jini lookup services, and the local lookup services will register themselves in the national level lookup services. Currently we are investigating the architecture of this distributed directory services system in order to make it robust, reliable and efficient. F.g. should the system be organized with the hierarchical structure like the domain name service (DNS) or the peer to peer structure like the Gnutella?<sup>11</sup> Each infrastructure element registers with attributes such as category, location, the server IP and URL, proxy interface program, and so on. Jini's attribute mechanisms support both type-based and content-based search styles, and make searching for particular attributes simple, quick, and effective. Standard taxonomies and specifications will describe infrastructures and their services accurately, but those taxonomies and specifications do not yet exist.

A typical system implemented with Jini has five basic concepts: Discovery, Lookup, Leasing, Remote Events and Transactions. Jini's ability to support spontaneously created, self-healing communities of distributed components is based on these concepts, and Java's Remote Method Invocation (RMI) and object serialization<sup>12</sup> techniques make the implementations of these concepts possible. Since the infrastructure web will be organized and integrated with Jini system, it will inherit these features from Jini. For example, with Jini's leasing concept, all infrastructures need to "sign" a lease with the lookup service in the registrations. Once the lease expires, the infrastructure will automatically be removed from the lookup service. In this way, a Jini lookup service has a self-healing ability for its directory management, and clients will not receive outdated information. The relationship between related infrastructures can be described using the remote event concept. For example, infrastructure #1 can tell a related infrastructure #2 which status messages it cares about, and infrastructure #2 will be automatically notified of changes by remote events from infrastructure #2. In this way, geographically distributed infrastructure elements can cooperate efficiently to detect, respond, and recover from the intrusion and attack. Using these Jini concepts, we believe that the infrastructure web can be easily implemented with the required characteristics that are proposed in section 1.

### **3.3 Sensor Web**

The ability to monitor the states of distributed infrastructure elements is another essential part of our system. In the analysis and detection of DDoS attacks, an analyst or upstream packet tracing system needs packet information from local machines as well as remote routers or firewalls. Here our sensor web system is actually a large-scale Distributed Smart Sensor Network (DSSN) that collects distributed sensor information both from intelligent software sensors as well as smart hardware sensors. The sensor web registers its sensors with the Directory Service, and the sensors provide distributed data sensing services to clients. Examples of sensors include computer network monitors (based for example on SNMP agents or packet analyzers), public health records, medical laboratory results, environmental monitoring stations, financial market trend monitors and so on. Human observations in the form of natural language reports are also relevant to this stage. Like the infrastructure servers, every sensor web system has an XML-web and a database server included. The XML-web is used to describe detailed sensor information and report real time sensor data on the web pages. After the sensor web system acquires the data from its sensors, it writes the data into its database periodically. All other remote applications or services can use ODBC or JDBC to access the data.

The advances in measurement devices have reduced cost to the point where it is now viable to develop large-scale distributed sensing systems. Meanwhile the advances in processor technology allow for relatively low-cost, low power, compact distributed processing integrated within these sensor devices, commonly referred to as smart sensors. Intelligent or smart sensors capable of parsing and filtering only the necessary or desired information allow for efficient use of memory, precious wireless bandwidth, and battery power needed for the transfer of sensor information. Before sending sensor information to related infrastructure elements, the sensor web system will preprocess the data from the distributed sensors using methods such as data filtering, data fusion and data mining. More information about our distributed sensor web systems can be found in.<sup>12</sup>

### **3.4 Emergency Information Search**

When an infrastructure element fails or is attacked intentionally, the damage needs to be assessed for a rapid recovery. In some cases, lost services might be covered by nearby or related infrastructure elements. The status of those other elements should be checked to help the coordinators to make informed decisions. Unfortunately, this kind of online emergency information search and response system does not exist at this time. The infrastructure web would have a nation-wide directory service of all critical infrastructure elements and could therefore perform the role of emergency response system.

Just like the “911” telephone emergency systems, the emergency information server should have a special and well-known domain name, and the emergency query forms should be well formatted. After clients submit the query, the HTTP server will transfer the query data to CGI or Java Servlet programs. These programs will process the query data and submit a formatted attributes template to the Jini lookup services. Then Jini systems will search through lookup services and return all relevant infrastructure element information and URL’s. Then, using the XML-based web, clients can check the real time status and internal states of other infrastructure elements that might assist in recovery.

## **4. RELATED WORK**

### **4.1. Grids Projects**

While commercial companies are making enterprise component-based frameworks available on the market such as Microsoft’s DCOM (Distributed Component Model) and Sun’s EJB (Enterprise JavaBeans), some more ambitious grids projects like Globus and CoABS are under the way for distributed computing. The Globus project is developing the fundamental technology that is needed to build computational grids. Grids are persistent environments that enable software applications to integrate instruments, displays, computational and information resources that are managed by diverse organizations in widespread locations. Meanwhile agents are a new software technology perfectly situated to take advantage of today’s computing infrastructure and software development. The DARPA CoABS research community is developing a prototype “agent grid” as an infrastructure for the run-time integration of heterogeneous multi-agent and legacy systems.

While these grids projects are focusing on how to build computational grids, our infrastructure web project intends to build specific grids for distributed infrastructure protection. Because all these grids projects try to use widespread existing computational and informational resources to support dramatically new classes of applications, the concern about framework, architecture and communication protocol for the projects should be similar.

### **4.2. UDDI**

The UDDI (Universal Description, Discovery and Integration) project is a commercial effort that shares some of the same objectives of the infrastructure web.<sup>14</sup> The UDDI project is an open industry initiative enabling businesses to discover each other and define how they interact via the Internet and share information in a global registry architecture. UDDI aims to be a mechanism that will enable businesses to find and transact with one another via their preferred applications. UDDI is also a framework for Web services integration that contains specifications for service description and discovery. The UDDI initiative is currently led by three companies (Ariba, IBM, and Microsoft), but eventually it will be turned over to a standards organization. The UDDI Business Registry is operated as a Web service supporting the UDDI specifications, and an Operator's Council helps set policy and quality of service issues for operators.

The UDDI vision of a global registry is similar to the infrastructure web plan, and we might leverage UDDI technology and/or standards to realize the infrastructure web vision. In the abstract, the plans are similar in that both are methods for connecting disparate elements offering a wide variety of services via flexible lookup services. Where commercial companies seek to conduct business transactions, infrastructure elements will share critical information. Certainly, the standardization effort associated with UDDI can serve as a guide for the difficult task of standardizing the format of published data from independent data providers.

## **5. SUMMARY**

The information revolution has introduced computers and the Internet into every corner of our society. Today we are relying more and more on the computer-controlled systems, but these systems are vulnerable to intrusion and destruction. The recent DDoS attacks against e-commerce companies have raised concerns about how to cope with cybercrime and cyberterrorism. Potential future attacks on critical infrastructure could cripple the nation by denying or disrupting the delivery of critical services. So how can we protect our national critical infrastructures from cyberterrorism? In this paper, we proposed six stages for the information infrastructure protection: intelligence gathering, analysis, interdiction, detection, response and recovery and our vision for how these stages can be integrated, with some detailed discussion on our plans for realization of an infrastructure web. The system will be a platform for decentralized monitoring and managing critical national infrastructures.

We believe that after the system is implemented and deployed, it will function as follows: while sensors web systems collect data for all critical infrastructures, the services on every infrastructure server monitor these data to verify that the infrastructure is running well. Once the infrastructure server detects attacks or failures, related infrastructure elements will be notified, and response and recovery steps will be taken automatically. Meanwhile, an emergency response coordinator can search and check real-time data from all appropriate infrastructures and choose the appropriate pre-simulated control strategy to respond to and recover from the attacks or failures. Success of this vision will depend on more than technology. The integration of the various agencies and organizations will require more complete knowledge of the operations of the various infrastructure categories from different fields. Organizations and domain experts will need to analyze the specific requirements for their areas, and more academic and industrial research on infrastructure protection is needed. Eventually state and local government cooperation and support will also be required to deploy this system.

### ACKNOWLEDGEMENTS

This research was supported by the Office of Science and Technology in National Institute of Justice (contract 2000-DT-CX-K001), the National Science Foundation (NSF contract CCR-9813744), Defense Advanced Research Projects Agency (DARPA contract F30602-98-2-0107) and the Air Force Office of Scientific Research (AfoSR contract F49620-97-1-03821). All opinions expressed here are solely those of the authors.

### REFERENCES

1. [http://www.sans.org/ddos\\_roadmap.htm](http://www.sans.org/ddos_roadmap.htm)
2. <http://www.cert.org>
3. <http://www.darpa.mil/iso/ia/>
4. <http://informant.dartmouth.edu>.
5. <http://www.ciao.gov>.
6. W.K Edwards, *Core Jini*, Prentice Hall, 1999.
7. I. Foster and C. Kesselman, *The Grid Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, San Francisco, 1998.
8. <http://coabs.globalinfotek.com>.
9. E.R Harold, *XML Bible*, IDG Books Worldwide, 1999.
10. D.P.Bertsekas, J.N.Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, MA, 1996.
11. <http://www.gnutella.wego.com>.
12. J.L.Weber, *Using Java 1.2*, Que Publishing, 1998.
13. G.C. Michael, C. Okino. *A Study of Distributed Smart Sensor Networks*, Dartmouth College, Thayer School of Engineering, Technical Report Preprint, March 2000.
14. <http://www.uddi.org/>.