

The DETER Project

Advancing the Science of Cyber Security Experimentation and Test

Jelena Mirkovic, Terry V. Benzel,
Ted Faber, Robert Braden, John T. Wroclawski
USC Information Sciences Institute
Marina Del Rey, CA, USA
{sunshine, tbenzel, faber, braden, jtw}@isi.edu

Stephen Schwab
Sparta, Inc.
Columbia, MD, USA
Stephen.Schwab@cobham.com

Abstract — Since 2004, the DETER Cybersecurity Testbed Project has worked to create the necessary infrastructure—facilities, tools, and processes—to provide a national resource for experimentation in cyber security. The next generation of DETER envisions several conceptual advances in testbed design and experimental research methodology, targeting improved experimental validity, enhanced usability, and increased size, complexity, and diversity of experiments. This paper outlines the DETER project’s status and current R&D directions.

Keywords: *cyber-security, testbed, experimental research*

I. INTRODUCTION

Cyber systems have become an inseparable part of our everyday lives. During past decade the Internet has permeated business and leisure sectors. More recently critical infrastructures are beginning to shift from physical to cyber control, and in some cases interconnecting with public data networks, to gain function, efficiency and cost-effectiveness. At the same time cyber attacks continue to increase in frequency and severity. These emerging trends – the increase in cyber attacks and the increase in potential damage these may create in our everyday lives – make cyber defenses a top funding priority for US government and industry.

Over the past 10 years heavy R&D funding of cyber security technologies has produced many promising approaches. However, large-scale deployment of security technology sufficient to protect vital infrastructure is both resource-intensive and often lacking. More importantly, most responses remain reactive, responding to individual new threats and attacks with patches and system modifications, rather than focusing on proactive design of more robust architectures and technologies.

We argue that an important contributor to this deficiency is the lack of experimental infrastructure and effective scientific methodologies for developing and testing next-generation cyber security technology. Specifically, current impediments

This research was sponsored by the US Department of Homeland Security and the Space and Naval Warfare Systems Center, San Diego (contract number N66001-07-C2001), and the US National Science Foundation (contract number CNS-0831491). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of Homeland Security or the Space and Naval Warfare Systems Center, San Diego.

to evaluating networked system security mechanisms include lack of support for effective research methodologies; inadequate models of attack and defense mechanisms; inadequate models of key inputs such as network structure, background traffic and mission traffic; and lack of testbed support for large-scale experimentation.

To address these shortcomings, the US Department of Homeland Security and the US National Science Foundation have funded the DETER project since 2004. DETER, led by USC/ISI, UC Berkeley and Sparta, Inc. is creating the necessary infrastructure—facilities, tools, and processes—to provide a national resource for experimentation in cyber security.

The next generation of DETER envisions a major advance in testbed design, focusing on dramatically improved experiment validity, together with significantly enhanced usability and increased size, complexity, and scope of experiments. The advanced conceptual design for next-generation DETER will stretch the envelope of testbed technologies and methodologies to influence the direction of future network testbed design. It will also provide an increasingly sophisticated and capable experimental platform as a shared service for a growing set of users, enabling world-class cyber security research.

This paper briefly describes the DETER testbed, and presents the project’s current research directions aimed at advancing the science of cyber security experimentation.

II. THE DETER PROJECT

A. What is DETER?

The DETER project consists of an operational *testbed facility* centered on experimental cyber-security research, test, and evaluation, together with a *research program* aimed at substantially advancing the infrastructures, tools and methodologies that underlie this experimentation. The DETER testbed [1][2][3] is hosted at the University of Southern California’s Information Sciences Institute and at University of California at Berkeley. The facility currently includes approximately 400 general-purpose computers and 10 FPGA-based reconfigurable hardware elements, richly interconnected

by a dynamically reconfigurable, switched network. Resources of the testbed are controlled and allocated to individual experiments using an extended version of the *Emulab* [4] testbed control software. Users receive exclusive, hardware-level access to the number of machines they need, and may set up network topologies, operating systems, and applications of their choice.

DETER is an open testbed, developed and supported as a national resource by DHS, NSF, and additional sponsors. DETER is administered using a project model. Any researcher is eligible to become a project member. Eligibility of project leaders is verified by confirming that the applicant is employed by a known institution or otherwise has roots in the community, for example through the production of prior research output such as whitepapers, papers, presentations or product catalogs. Project leaders approve membership of regular users under their project, thus e.g., a faculty member would approve his/her students. There is no cost to use the DETER testbed and tools. There are also no special technical requirements. The testbed can be accessed from any machine that runs a web browser and has an SSH client. Experimental nodes are accessed through a single *portal* node via SSH. Under normal circumstances, no traffic is allowed to leave or enter an experiment except via this SSH tunnel.

B. Who Uses DETER?

The DETER testbed and tools have been used by 1387 users (as of August 2010), from 14 countries and 4 continents. The majority of DETER users are located in the North America and Europe, but lately we are seeing an increased use in regions such as Middle and Far East, India and South America. Around 70% of our users come from academia, 20% from industry and the remainder from government schools and organizations.

C. What Can DETER Do Today?

DETER is used today to support a variety of cyber-security research. The largest use occurs in the following research fields: comprehensive security systems, privacy, spoofing, worms, product evaluation, overlays as security solutions, attack traceback, multicast security, wireless security, intrusions, botnets, denial-of-service, malware, infrastructure security and new security architectures. A significant number of users use DETER as a platform to test new testbed technologies, or to learn how to deploy facilities similar to DETER at their own institutions. Finally, in the last three years we have seen a dramatic increase in number of projects and users that use DETER in classroom environments, to supplement regular cyber-security education with practical exercises. These exercises increase student interest in material and improve comprehension of topics taught in class. The project strongly supports such use. DETER is currently setting up a Moodle [5] server, slated to be made public in Fall 2010, that will host educational content and facilitate its sharing between teachers, and is adding account management and

scheduling enhancements designed to simplify classroom and educational user administration.

The project has developed a number of new capabilities for security research, in addition to the range of useful abilities inherited from the Emulab software. Among those deployed today are:

- An integrated experiment management and control environment called SEER [6][7], with a rich set of traffic generators and monitoring tools.
- The ability to run a small set of risky experiments in a tightly controlled environment that maximizes research utility and minimizes risk [8].
- The ability to run large-scale experiments through federation [9] with other testbeds that run Emulab software, and more recently with facilities that utilize other classes of control software.

D. What Have We Learned?

While DETER's contributions in the five years of its existence have been significant, the experience of running the testbed and developing tools for cyber-security experimentation has helped us recognize significant new challenges in this field. We believe that these challenges are major obstacles to easy, versatile, realistic, scientific and repeatable experimentation needed to advance the state of cyber defense research. These challenges range from those that are purely technical to those that are primarily human-oriented and community-focused.

A first challenge lies in the *diversity of users and uses* of network testbeds, and how to best support all at a fixed budget and within a shared infrastructure. Many testbed users have insufficient systems knowledge to effectively use the facility, while others may benefit from a focus on increased scientific rigor. Each of these classes of users benefit from significant help to learn how to most effectively use the testbed for their research.

On the other hand, there are a few sophisticated users that require advanced capabilities and need little guidance to produce highly visible results in terms of publications and new products. Here, the challenge is to most effectively get out of the way, ensuring that mechanisms established for ease of use and guidance of mainstream users are not impediments to use by this sophisticated user class.

Finally, there is a large diversity of testbed use patterns between the research, industry/government and education communities. Researchers tend to investigate new phenomena in a collaborative manner, and build unique tools for each new project they undertake. Industry/government users tend to focus on system evaluation and expect a standardized evaluation environment and test cases. Educational users need a special set of protections to minimize cheating and manage both desired and undesired sharing of work across groups. They further may lack sufficient systems background to

Design ↔ Instantiation ↔ Execution ↔ Analysis




Figure 1: Experiment lifecycle

intuitively “jump in” to using a testbed, which leads to a steep learning curve and few useful results.

A second challenge lies in the *fast pace with which the cyber-security field is moving*. As new technologies, attacks and defenses appear with great frequency, it is difficult for testbed development efforts to offer detailed support for experimentation in all emerging areas. Creation of new tools and mechanisms requires significant domain expertise and time investment. By the time support is available it may already be obsolete. We believe that the only sustainable approach is to adopt an overall focus on proactive development of broad capabilities, rather than reactive response to specific, narrow risks and attacks. We combine this proactive overall focus with an effort to foster and catalyze domain-specific user communities, offering motivation and support for these communities to share, exchange, and reuse the latest information and tools.

A third challenge lies in the need for determining and then creating appropriate *realism* in the experimental environment. Two separate factors are important. First, for any particular experiment, it is necessary to understand which aspects of the experiment must accurately reflect the technology’s actual intended operating environment (the “real world”), and to what degree. Second, for each such aspect, it is necessary to develop an appropriate model or configuration for the experiment.

For example, consider an Internet distributed denial of service (DDOS) experiment. There is very little detailed data available about Internet topology and traffic, and what is available must be heavily transformed to create experiment-specific models appropriate for use in a smaller-scale, low-diversity testbed. This adds significant time and challenge to experimentation that may not directly lead to publishable research, but is necessary to facilitate it. Users tend to view this as overhead, and consequently often design naïve experiments that minimize this overhead but lead to flawed methodologies or results.

A fourth challenge lies in the *unpredictability and complexity* of working with real hardware and software, especially at

large scale. It is difficult for users to verify that their experiment’s behavior matches what is intended, and that the results contain no misleading artifacts. It is even more difficult to diagnose the source of experimental artifact and error when it occurs.

III. NEW DIRECTIONS FOR CYBER-SECURITY RESEARCH AND DETER

The DETER project’s vision of the future of cyber security experimentation is as large a step forward from the current state of the art as the current DETER testbed is from the small isolated testbeds traditionally seen before DETER. The future DETER testbed will:

- support larger and more complex experiments,
- advance the scientific quality and accuracy of experimental results,
- provide adaptive and domain-specific tools to help users create new experiments and to deal with the great complexity of large-scale cyber security experiments,
- build a knowledge base of experimental designs and results,
- incorporate an extensible software architecture,
- provide a user-friendly interface for both novice and experienced users,
- and, as a result of these advancements, support a significantly larger and more diverse research community.

We now provide more details about three major research and development initiatives the DETER project is undertaking to reach this future: advancing the *science* of experimentation, advancing *testbed technology*, and supporting new *application domains*. While DETER focuses on cyber-security experimentation, these initiatives will advance the state of network testbeds in general.

A. Creating an Advanced Scientific Instrument

As a scientific instrument, a network testbed must provide repeatability, validity, and usability. It should advance the scientific enterprise by helping experimenters to distinguish valid results from artifacts, and to build on each other's work. Furthermore, it should provide a significant expansion of experiment scope beyond that available today, for example: support for structured experiment design and worst-case experiments, multi-party experiments, multi-dimensional experiments, and experiments that interact safely with the outside world. We briefly outline each of these areas:

- a. *Worst-case experiments* focus on explicitly identifying and exercising worst-case or corner-case behaviors, pushing to the breaking point one or more dimensions of the technology or scenario under investigation. These experiments closely resemble attack scenarios from real life. Yet they are fundamentally different from experiments commonly carried out today in network testbeds, which avoid extremes to ensure determinism and to focus on behavior when the technology is operating properly. They are particularly challenging in cases where the worst-case evaluation may misleadingly trigger limitations of the testbed infrastructure, masking the behavior of the actual scenario under test.
- b. *Multi-party experiments* involve experimental scenarios with two or more participating entities (human or algorithmic), each with only partial visibility into the modeled universe. These experiments closely resemble security interactions in the real world. For example, a multi-party experiment might combine experimenters playing different roles (e.g., attacker/defender or local facility administrator/global network security operator), or experimenters with different levels of expertise. Consider a concrete scenario. In an experiment involving a malware defense system, the system and its operators may have detailed knowledge of their own network and host environment, but only limited and inaccurate information about the overall network topology and the actions of the distant attacker. Actions taken to observe, analyze, and undermine the attack must be carried out through this lens. Current network testbeds do not support such information hiding and fine-grain visibility control.
- c. *Multi-dimensional experiments* investigate multiple dimensions (technologies, phenomena, etc.) simultaneously. Such experiments are difficult to specify and construct, yet often crucial to obtaining useful results. It is of limited use to investigate individual technologies (OS, application, user-interface, node security, network routing, network transport, network authentication, network firewall, network IDS, etc.) in isolation, due to coupling effects and system-level behavioral constraints. One complexity of security experiments is created by the reality that security problems and properties are emergent in the collection

of technical elements (e.g. inherently complex), and it is rarely the case that simpler scenarios can simply be studied independently and then composed or "added together".

- d. Experiments that can *interact with the larger world outside the testbed* are necessary to study current attack trends. For example, many useful experiments require some degree of interaction with the Internet, to capture Internet properties such as fidelity, scale, and non-determinism, or to interact with malware "in the wild". Yet such experiments are potentially risky. Even if built for benign purpose, these experiments may lead to unexpected and undesired interaction with the outside world, such as accidental use of the testbed to spread malware to outside systems. What is required is, first, methodologies to analyze and reason about such risks, and second, new testbed technologies that preserve the properties essential to the experiment but rigorously minimize the risk.

We believe that the fundamental change needed to address the above issues is a shift from focusing only on the runtime configuration of an experiment at a highly concrete level, as is the current practice in network testbeds, to capturing and describing an experiment's entire lifecycle at a more abstract level, as shown in Figure 1. This includes such elements as the experiment definition (topology, node configuration, traffic generators, event generators, monitors), the workflow (set of events that should occur in order or with specific timings to define the experiment), the invariants (conditions necessary for experiment success) and the analysis (types of statistics collected and how they are processed and visualized). This level of experiment description will advance the value of DETER as a scientific instrument by:

- Supporting the use of models for classes of experiments. A *model* is an abstract representation of an entire class of experiments such as "worm propagation experiments". It defines actors (e.g., vulnerable and infected hosts), their actions (e.g., once infected, start scanning randomly), experiment invariants and relevant monitoring and analysis.
- Permitting the refinement of models into complete recipes. This refinement occurs through concrete realization of a model, including an interaction with users to bind specific values to parameters (e.g., type of scanning = random). The final outcome, a recipe, is a complete specification of an experiment that can be directly run on the testbed.
- Supporting replay and automatic operation of experiments, including "clusters" of experiments that explore different points in a design space. This support might include such capabilities as
 - Automatically invoking instrumentation and configuring visualization of results.
 - Facilitating the standardization of experimental conditions, and facilitating the execution of multiple

experiments while varying key parameters or conditions.

- Facilitating the sharing of experimental setups and results as recipes or models.

This shift towards a richer model of experiment definition represents a change in the basic research paradigm. It will enable researchers to focus on their research and to create sophisticated experiments that properly evaluate performance of proposed technologies and algorithms. Better evaluation practices and increased sharing should help to transform research practice in the cyber security community from single-point, naïve efforts into collaborative, evolving, and sophisticated endeavors.

Another consequence of this shift will be to enable users to build shared repositories of tools, data, experiments and models for specific, popular research problems. This is very important from the DETER project viewpoint, as it will foster a more sustainable model of open-source and community collaboration instead of the current centralized development model, in which DETER staff provides tools and help to all users.

Working towards this overall paradigm, a number of detailed steps may be considered. The following paragraphs discuss a subset of these details, which represent near-term aspects of our work towards achieving the overall goal of creating an advanced scientific instrument.

An experiment is defined along a number of *dimensions*. Some dimensions are well known – e.g., topology or traffic characteristics – while other dimensions may be specific to a specific type of experiment (e.g. type of scanning for a worm spread experiment). Dimension descriptions may range from very concrete (as in a recipe) to very abstract (as in a model). A model may generate detailed values for its dimensions, or it may be queried to provide the details when needed. A model

may represent only a subset of the possible dimensions. Composition of models for different dimensions will provide a natural way to construct, reason about, and manipulate large, complex experiments.

A major hindrance to effective testbed use is the *steep learning curve* needed to set up and orchestrate experiments. Ideally, it would be possible to create a powerful but user-friendly experimenter support system that meets the needs of different levels of users – from novice to sophisticated – and all user activities – from classroom exercises to product testing to scientific research. The interface framework for designing and executing experiments should be unified, flexible, and easily extensible. We use the term “*workbench*” for such an experimenter interface. This term pertains to a general user interface and the backend that helps users design and manipulate experiments.

Risky experiment management is enabled by applying a wide variety of techniques, from a priori constraints placed on an experiment's implementation, to verification of key assumptions, to external enforcement of invariants by the testbed infrastructure. All of these – constraint, verification, and enforcement – rely upon formal representations of the assumptions and behaviors that describe the semantics of the experiment, and reasoning about these assumptions and behaviors to create constraints that preserve function but manage risk. These constraints are divided between those defined by the experimenter – who knows the precise goals and objectives of the experiment – and those defined by the DETER testbed operators – who know the overall assurance goals of the testbed, and the potential vulnerabilities and mitigations available through different enforcement mechanisms [8].

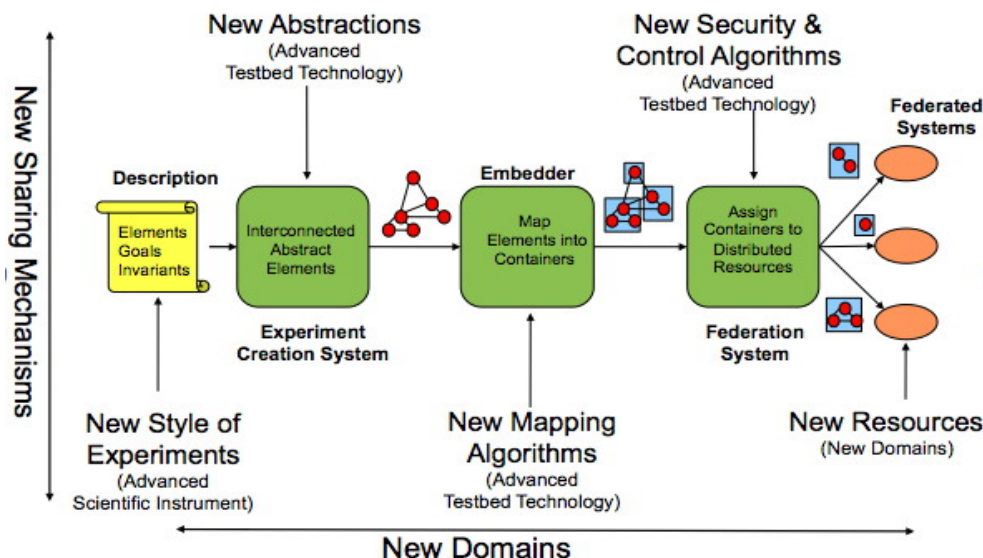


Figure 2: Virtualization and embedding process in future DETER testbed

B. Advanced Testbed Technology

To explain this new capability, we contrast it with DETER's current Emulab-based mechanisms. Today, an experiment description contains little more than the network wiring diagram and some initial configuration information for each node. An Emulab mechanism operating within DETER maps this description to a set of physical nodes running the selected OS and interconnects these nodes with VLANs that match the network diagram.

In contrast, the advanced testbed technology we are developing takes as input the experiment description created by a researcher, using the advanced scientific instrument mechanisms described in the previous section. Experiment descriptions are composed of elements that may range from concrete – physical nodes loaded with specific OS and application software – to abstract models that can be virtualized in a number of *different* ways, depending on the required fidelity and experiment objectives. To guide this process, descriptions incorporate semantic information such as goals and invariants associated with the experiment. A novel translation step, the *virtualization engine*, will then examine each element and determine the appropriate physical or virtual technology to realize that element on the testbed. Then, a generalized *embedder* will allocate and configure the testbed hardware and software resources needed to instantiate this set of elements. Lastly, the *federator* will allocate the embedded containers to physical nodes – optionally using remote as well as local resources. The virtualization and embedding process is shown in Figure 2.

We now give some examples of elements and the resource containers in which they may be embedded to illustrate this novel testbed technology. An element may be a single computer running an OS, and as in Emulab, embedded as a physical node running that OS, or alternatively as a virtual guest OS on a virtualization layer that provides the necessary properties to satisfy all invariants, including invariants for performance (validity management) and containment of malicious behavior (risky experiment management). On the other hand, an element could be substantially more abstract – for example, a host node within a 200,000 node botnet – that would be embedded as thousands of 'bot emulation threads on hundreds of physical machines. An element could also be a cluster of nodes together emulating the routing, loss, and delay properties of a core network. Finally, an element could be a specific switch, router, or firewall device – but rather than using the real device or emulating the device on a general-purpose PC, DETER would implement the device by loading and configuring an FPGA-based platform – achieving both high performance and a high degree of fidelity for complex parallel algorithms – if that was the requirement for a given experiment.

In summary, the result of deploying an experiment on the future DETER testbed will be an engineered environment, precisely tailored to support this experiment. It will consist of an interconnected set of containers, each of which may be a

physical node, a virtual machine, a simulator thread realizing some experiment-specific model, an FPGA-based platform with customized firmware, or some other environment. The virtualization and embedding process will not be restricted to a fixed set of container classes nor to a fixed set of federated testbed types. New containers classes – virtual machine technologies, simulators, emulators, or hybrids – will be added as plug-ins to the virtualization engine to allow evolution of the models we can support. Similarly, the current DETER federation architecture [10] will be augmented with additional plug-ins to support the use of testbed resources – both hardware and software – on remote testbeds using both Emulab-derived and radically different testbed control software – such as a custom SCADA testbed.

This envisioned capability subsumes all of the current functionality of the present DETER testbed and DETER federation architecture, while radically generalizing in three dimensions. First, it enables the target containers to be not only physical nodes, but also any virtualization technology in the most general sense – any configuration of hardware and/or software that we engineer to re-create the behavior of an element of a cyber-experiment *at an appropriate level of abstraction for that experiment* – that is, with the semantics needed for experimental validity in the case being addressed. Second, it enables the creation of multi-party experiments, above and beyond the use of federation to exploit resources on remote testbed facilities. This is supported through an ability to combine the recipes for different experiments into large composed experiments, while respecting the semantics of each sub-experiment including hiding or controlling information flows between elements of the federated experiment. Third, it provides for the mapping between the scientific recipe describing the experiment and the engineered containers realizing the experiment to be informed and controlled by the semantics of the experiment – so that different mappings are possible depending on exactly what the goals and objectives of the researcher are for a given experiment.

C. New Application Domains

The goal of any network testbed, including DETER, is to support research in key and emerging areas. While it is difficult to predict which security areas will dominate the field in the future, we briefly describe here the support we plan for three specific domains that currently satisfy this criteria: large-scale semi-self-organizing systems, critical infrastructure, and wireless.

1) Large-Scale Semi-Self-Organizing Systems

Botnets are a current example of large-scale, semi-self-organizing systems (hereafter, SSOs) that represent powerful and versatile platforms for attackers. These systems are ultimately characterized by the overall aggregate behavior of thousands or millions of elements, rather than the individual behavior of a single element.

Cyber security researchers will study this class of threat for years to come. SSOs pose a fundamentally new domain of

malware whose systematic investigation will require a suite of very different tools and techniques in DETER. SSO's differ from viruses, worms and earlier malware in that the system operates semi-autonomously, yet the attacker is actively engaged at some level, managing the system, uploading new software modules, and monitoring its effectiveness, anti-reverse engineering protections, and rate of growth and persistence on the Internet. In short, SSO's can be thought of as exhibiting a form of guided intelligence. The current generation of botnet experimentation tools, such as SLINGbot [11] and Rubot [12], emulate only the botnet communication protocols. Emulating guided intelligence behavior at large scale, or combining simulation with testbed emulation, is needed to obtain realistic experiments in this research field.

2) Critical Infrastructure Support

The intersection of cyber security with critical infrastructure lies in *cyber-physical systems*. Such systems, including power grids, water, oil and gas distribution systems, control systems for refineries and reactors, and transportation systems are all vulnerable to attacks in both the cyber and physical realms, separately or, more dangerously, in combination. Protecting such systems requires the ability to model the reaction of such cyber-physical systems to both kinds of attacks, and, more interestingly, modeling the coupled behavior of the cyber and physical systems to a single attack.

Our objective is to extend DETER to provide useful modeling of such systems. Cyber-physical systems can be modeled in DETER by simulating the physical system on some DETER nodes operating as cluster computers, while emulating or modeling the cyber control system using standard DETER capabilities. Ties between the cyber and physical elements, corresponding to the sensors and effectors that are present in such infrastructure, must be modeled or emulated as well.

Critical infrastructure systems are naturally federated, with portions of the physical infrastructure owned and operated by multiple utilities, operating under oversight by multiple jurisdictions. In the longer term we envision the ability to model individual domains within such infrastructure on separate testbeds, and to use testbed support for federation to allow for participation in the modeling by actual operators, with support for information hiding. For example, utilities could participate in emulations of the larger system, modeling their own response, without compromising the own sensitive information on which those responses are based.

3) Wireless Support

Wireless security is an important and complex research field that is not presently well supported by DETER. An ultimate goal would be to enable long-term research covering current and future wireless technologies as well as local area, cellular-based, and wide area RF communications, and to support network applications on many different platforms such as PCs, handhelds/PDAs, and sensor networks. By adding wireless support within the existing DETER security experimentation

framework we can leverage experimental tools such as traffic generators that are common to all types of security research, while integrating emulators and wireless-protocol specific tools that span the physical link, network infrastructure, and wireless application domains.

A crucial restriction with respect to effective support of wireless experimentation is DETER's current reliance on the "nodes and links" resource allocation model inherited from Emulab. Wireless communication, in contrast, requires a different underlying model more tuned to broadcast media and communication spaces. Development and implementation of appropriate underlying semantics for wireless support at the infrastructure level is a major challenge.

IV. CONCLUSIONS

The cyber threat changes rapidly, with an evolving landscape of cyber attacks and vulnerabilities that require technologies for prevention, detection, and response. One key aspect of a cyber security strategy is the scientific test and analysis of new and emerging technologies for cyber space. The DETER cyber security testbed provides a unique resource for such testing, and the work described here aims to provide new science-based methods, tools, and technologies to advance the field of cyber experimentation and test. In addition to the technology-based contributions we are also engaged in efforts to increase the community of testbed users, educate future researchers and the next generation of computer operations staff, and provide a platform for training of cyber security personnel and postmortem analysis of their actions.

REFERENCES

- [1] DETER Testbed Web page, <http://www.deterlab.net>
- [2] DETER Project Web page, <http://www.isi.edu/deter>
- [3] Terry Benzel, Robert Braden, Dongho Kim, Clifford Neuman, Anthony Joseph, Keith Sklower, Ron Ostrenga and Stephen Schwab, "Experience with DETER: A Testbed for Security Research," in *Proceedings of Tridentcom*, March 2006.
- [4] Emulab Testbed Web page, <http://www.emulab.net>
- [5] Moodle Course Management Tool, <http://www.moodle.org>
- [6] SEER Experimentation Workbench, <http://seer.isi.deterlab.net>
- [7] Stephen Schwab, Brett Wilson, Calvin Ko, and Alefiya Hussain, "SEER: A Security Experimentation Environment for DETER," in *Proceedings of the DETER Community Workshop on Cyber Security Experimentation and Test*, August 2007.
- [8] J. Wroclawski, J. Mirkovic, T. Faber and S. Schwab, "A Two-Constraint Approach to Risky Cybersecurity Experiment Management," Invited paper at the *Sarnoff Symposium*, April 2008.
- [9] T. Faber and J. Wroclawski, "A Federated Experiment Environment for Emulab-based Testbeds," in *Proceedings of Tridentcom*, 2009.
- [10] DETER Federation Architecture, <http://fedd.isi.deterlab.net>
- [11] Alden W. Jackson, David Lapsley, Christine Jones, Mudge Zatkan, Chaos Golubitsky, W. Timothy Strayet, "SLINGbot: A System for Live Investigation of Next Generation Botnets," in *Proceedings of the 2009 Cybersecurity Applications & Technology Conference for Homeland Security*, pp. 313-318, 2009.
- [12] Christopher Patrick Lee, "Framework for Botnet Emulation and Analysis," Ph.D. Thesis, Georgia Institute of Technology, March 2009.