

Cyber-Physical Testbeds: Scientific Instruments for Cyber Security Assessment of Critical Infrastructures

Christos Siaterlis and Béla Genge*
(*contact author)

1. INTRODUCTION

Modern societies depend to a large degree on the quality and reliability of the services that Networked Critical Infrastructures (NCIs) provide. Physical infrastructures, such as transportation systems, the electricity grid, and telecommunication networks, provide fundamental services for the smooth functioning of the economy and for the lives of citizens. Therefore, accidental or intentional failures of these infrastructures represent one of the most important risks that our society faces.

During the last years we have witnessed a dramatic increase in the use of Information and Communication Technologies (ICT) within such NCIs. The motivation was mainly to reduce the costs of industrial installations and to implement new services such as the remote monitoring and maintenance of infrastructures, energy markets, and the emerging smart grid. Although the advantages of this trend are indisputable, the downside is that widespread use of standard ICT components exposes these vital NCIs to significant but nonetheless common cyber threats. For instance, deliberate attacks through computer malware [6] or unintentional threats from misconfiguration and software bugs within ICT systems [5] can lead to severe service outages.

This fact has also been highlighted by several studies and reports concerning the security of Supervisory Control And Data Acquisition (SCADA) systems [6, 15]. SCADA systems represent the core infrastructure of NCIs, providing the capabilities for monitoring and controlling of physical processes. They mainly consist of actuators, sensors and hardware devices that perform a physical action, e.g., open a valve, as well as all the ICT devices and software that monitor and control physical processes. Unlike traditional ICT systems where the effects of disruptive cyber attacks are mostly limited to the cyber realm, in the context of critical infrastructure assets, such attacks can result the loss of vital services such as transportation, water and gas supply. To properly assess the impact of cyber threats against both the physical and cyber dimension of NCIs an accurate and efficient scientific instrument for conducting experimental tests and measurements is needed.

Cyber-physical testbeds that actively support the “scientific method” are a clear example of such *scientific in-*

struments. Testbeds may be developed leveraging real systems, emulators or software simulators. Unfortunately, experimentation with production systems for security and resilience tests entail risks of potential side effects to mission critical services [10]. Similarly, the development of a dedicated experimentation infrastructure with real components in order to run disruptive experiments involves safety risks [10] and has high installation costs [11]. Software based simulation has always been considered an efficient approach to study some physical systems, mainly because it can offer low-cost, fast and accurate analysis. Nevertheless, it has limited applicability in the context of cyber security due to the diversity and complexity of computer networks. Software simulators can effectively model normal network conditions, but fail to capture the way computer networks fail [7]. On the other hand, in many cases, emulators can capture not only whether a system will fail but also how it will fail.

To address the existing need for cyber-physical testbeds, in this article we present a novel *Experimentation Platform for Internet Contingencies* (EPIC). EPIC is a modern scientific instrument that can provide accurate and repeatable assessments of the impact that cyber-attacks may have on the cyber and physical dimensions of NCIs. To model the complexity of today’s NCIs, EPIC uses a computer testbed based on Emulab [20, 24] to recreate the cyber elements of a NCI and software simulators for the physical components.

2. MOTIVATION

A major limitation of existing testbeds is the inability to run security experiments on multiple heterogeneous NCIs. In fact, today’s NCIs are highly interconnected and interdependent, which means that a single failure within one NCI might have a cascading effect on others. For example, the collapse of India’s northern electricity grid in July 2012 affected more than 600 million people and led to the loss of power in transportation, health care, and many other sectors. Scenarios such as this one (that might also be caused by cyber attacks) need to be recreated, analyzed and understood in a laboratory environment in order to develop the necessary security measures that can be applied in real settings.

By recreating key connections between the cyber and physical dimensions of NCIs, EPIC provides a diverse palette of research applications. Besides typical examples such as vulnerability testing, impact analysis, and validation of different techniques, EPIC can provide the necessary tools for closing an important loop in cyber-physical experimentation: human operators. In the context of NCIs, human op-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2008 ACM 0001-0782/08/0X00 ...\$5.00.

erators play a significant role in ensuring the stability and normal functioning of physical processes. Human operators can directly interact with EPIC as part of an experiment or they can be simulated by modeling their standard operating procedures. Either way, EPIC can be used in the future to build complex experiments, which may test the effect of commands issued by human operators on physical processes or measure the reaction of human operators to the changes in the state of physical processes. Consequently, we consider that EPIC brings an important development in today’s experimentation testbeds by providing more accurate experiments that are closer to the real operation of NCIs.

2.1 Testbed Requirements

A cyber-physical testbed as a modern scientific instrument needs to be compatible and actively support the “*scientific method*”. The instrument should actually enable researchers to apply rigorous scientific methods by ensuring the fidelity, repeatability, measurement accuracy, and safe execution of experiments [20].

2.1.1 Fidelity

Experimentation testbeds need to reproduce as accurately as possible the real system under study. However, in many cases reproducing in an absolute way all details of a real system might not be necessary. Therefore it is preferable for an experimental platform to offer an “adjustable level of realism”, meaning that we can use the level of detail that is sufficient to test the experiment hypothesis. For example one experiment might need to reproduce a network at the very low level using real routers while for another experiment the use of a software router might be sufficient. The concept of adjustable level of realism is to have the option to use real hardware when it’s really needed and emulators, simulators or other abstractions when not.

2.1.2 Repeatability

This requirement reflects the need to repeat an experiment and obtain the same or statistically consistent results. Repeatable experiments require a controlled environment, but to achieve them the researcher has to define clearly and in detail the experiment’s initial and final state as well as all events in between these two states. To reproduce a previously stored experiment scenario the researcher should be able to setup the experimental platform in the initial state and trigger all necessary events in the right order and time of occurrence.

2.1.3 Measurement Accuracy

Experiments should be accurately monitored and measurements should not interfere with the experiment in such a way that they might alter the experiment’s outcome. Therefore separation of control, measurement and experiment processes is needed.

2.1.4 Safe execution

In most cases security experiments assume the presence of an adversary that employs malicious software to reach his/her goals. The effect of this software can be unpredictable and may have disruptive effects on physical systems. Such cases need to be recreated in experiments, but without jeopardizing the physical testbed itself and without threatening the researchers.

2.2 Existing Approaches

In order to assess the state of the art, we performed a literature review and evaluated the features of currently available testbeds against the previously defined set of requirements. A summary is given in Table 1, where ‘●●●’ was used to denote a strong support for a specific feature, i.e., currently available, while ‘●’ denotes that the feature is weakly supported and in several cases it is currently unavailable, although it might become available in the future.

The US National SCADA TestBed (NSTB) program [23] constitutes a national collaborative laboratory project intended to support industry and government efforts to enhance the cyber security of industrial installations. It provides a wide range of facilities to recreate real-world systems from generation to transmission, which includes real power grid components as well as industry-specific software products. The NSTB was successfully used to identify vulnerabilities and to harden the protection mechanisms of control systems. Nevertheless, the cost of deploying a similar installation limits its practical applicability in multi-domain heterogeneous cyber-physical systems.

A collaborative effort between ENEL SPA, Italy, and the Joint Research Centre, Italy, led to the development of a protected environment recreating the physical characteristics of a real turbogas power plant [15]. The testbed accurately reproduces both the cyber and physical characteristics of a typical power plant. It includes a scaled down physical process, typical field networks, process network, demilitarized zones, horizontal services, corporate domain and real industrial standard software. The testbed has been used to analyze attack scenarios and to test countermeasures in a safe environment. Unfortunately, as previously stated, the high fidelity of pure physical testing environments is counterbalanced by their poor flexibility and the high cost of maintenance of similar architectures.

The cyber-DEFence Technology Experimental Research (DETER) testbed [2] is an Emulab-based testbed providing repeatable security-related experiments. DETER is part of the DETER Enabled Federated Testbeds (DEFT), which intends to interconnect geographically distributed testbeds to enable experimentation in the cyber-physical space. Within the DEFT consortium DETER has been interconnected [25] with Virtual Power System Testbed (VPST), a testbed developed by the University of Illinois [3]. VPST provides simulation capabilities of electricity grids through real-time simulators such as PowerWorld and extends DETER’s capabilities to enable experimentation with cyber-physical systems. The key difference between EPIC and DEFT is that EPIC provides a cost effective and scalable solution for experimenting with multi-domain heterogeneous physical processes (through its software simulators), while the efforts within DEFT seem to be more focused on a specific infrastructure, e.g., the power grid. Nevertheless, EPIC can also be viewed as complementary to the DEFT initiative since the software simulators developed for EPIC could be easily reused within DETER.

The PowerCyber testbed developed at Iowa State University [11] integrates SCADA-specific hardware and software with Real-Time Digital Simulators to simulate electrical grids. The testbed uses virtualization techniques to address issues related to scalability and costs, and the ISEAGE project developed at the Iowa State University to enable wide-area network emulation. The testbed also provides

Table 1: Comparison of testbed features and cost-effectiveness. We used '●●●' to denote a strong support, '●●' to denote a moderate support, and '●' for a weak support of a specific feature.

	Fidelity		Repeatability		Meas. Accuracy		Safety		Cost-Effective	Multi CI
	Cyber	Physical	Cyber	Physical	Cyber	Physical	Cyber	Physical		
<i>Real cyber, Real physical</i>										
NSTB[23]	●●●	●●●	●●	●	●●	●●	●●●	●	●	●
ENEL-JRC[15]	●●●	●●●	●●	●	●●	●●	●●●	●	●	●
<i>Real cyber, Sim physical</i>										
EPIC	●●●	●●	●●●	●●●	●●●	●●●	●●●	●●●	●●	●●●
DEFT[25] (DETER+VPST)	●●●	●●	●●●	●●●	●●●	●●●	●●●	●●●	●	●
PowerCyber[11]	●●●	●●	●●	●●●	●●	●●●	●●●	●●●	●●	●
<i>Real&Sim cyber, Real physical</i>										
TUB[8]	●●	●●●	●●	●	●●	●●	●●●	●	●	●
<i>Real&Sim cyber, Sim physical</i>										
VCSE[14]	●●	●●	●●	●●●	●●	●●●	●●●	●●●	●●	●●●
<i>Sim cyber, Sim physical</i>										
SCADASim[17]	●	●●	●●●	●●●	●●●	●●●	●●●	●●●	●●●	●●●
HLA[16]	●	●●	●●●	●●●	●●●	●●●	●●●	●●●	●●●	●●●

non-real-time simulation capabilities, primarily used for simulating larger systems and for performing state estimation and contingency analysis.

An approach that uses real components for the physical dimension and partly simulated ones for the cyber dimension comes from the Tsinghua University of Beijing (TUB), China [8]. The approach uses real SCADA control servers, the NS-2 network simulator, combined with real control hardware and field devices. The testbed was designed to assess the impact of several cyber attacks on the SCADA system, including packet forging, compromising of access control mechanisms, and compromise of SCADA servers. Although such a testbed would provide reliable experimental data, since almost everything is real, it would be hardly able to support tests on large infrastructures such as an entire electrical grid.

Sandia National Laboratory developed the Virtual Control System Environment (VCSE) [14] with the purpose of exploring vulnerabilities, training operators, and validating mitigation techniques. The testbed employs OPNET to integrate real devices with simulated networks and PowerWorld as the power system simulator. VCSE also incorporates Umbra, Sandia’s patented framework, which provides a centralized environment for monitoring and controlling multiple simulated components.

The SCADASim framework [17] developed at the Royal Melbourne Institute of Technology (RMIT) University provides a set of predefined modules for building SCADA simulations. The framework employs the OMNET++ discrete event simulation engine to recreate typical SCADA components and to provide an underlying inter-model communications layer. SCADASim supports integration with real devices through modules implementing industry standard protocols. The framework can be used to develop a wide range of SCADA simulations and to evaluate the impact of cyber attack scenarios on communications and on the normal functioning of physical processes.

Finally, we mention a “system-of-systems” approach, developed at the Swiss Federal Institute of Technology, Zurich [16]. The testbed uses an implementation of the High Level Architecture (HLA) simulation standard to provide a multi-

domain experimentation environment that interconnects simulators from different domains. The testbed was designed to provide support for exploring “what-if” scenarios in the context of complex interdependencies between critical infrastructures. Unfortunately, such testbeds might prove effective on interdependency studies, but as already mentioned, they fail to accurately recreate the cyber layer.

3. EPIC OVERVIEW

The architecture of EPIC involves the use of an emulation testbed based on the Emulab software [20, 24] in order to recreate the cyber part of NCIs, and the use of software simulation for the physical components.

By employing an emulation-based testbed we ensure strong fidelity, repeatability, measurement accuracy and safety of the cyber layer. This approach is well-established in the field of cyber security [2] and was chosen in order to overcome the major difficulties that rise while trying to simulate how ICT components behave under attacks or failures.

For the physical layer EPIC uses simulation, since this provides an efficient, safe and low-cost approach with fast and accurate analysis capabilities. Although it weakens the fidelity requirement, software simulation enables disruptive experiments on multiple heterogeneous physical processes. Furthermore, today we can find complex models of several physical systems in the literature. By integrating them in software simulators the behavior of real physical systems can be accurately reproduced. A clear example in this sense is the energy sector, where simulation has become so accurate and trusted that it is commonly used to aid decision making between transmission system operators.

3.1 Recreating Cyber Systems

The use of emulation testbeds is becoming more popular. One of the most advanced software suites in this direction is Emulab [24]. Nowadays the software is actively supported by multiple universities and there are many private installations throughout the world.

We have developed in our laboratory a testbed using the Emulab architecture and software (Figure 1(a)). By adopt-

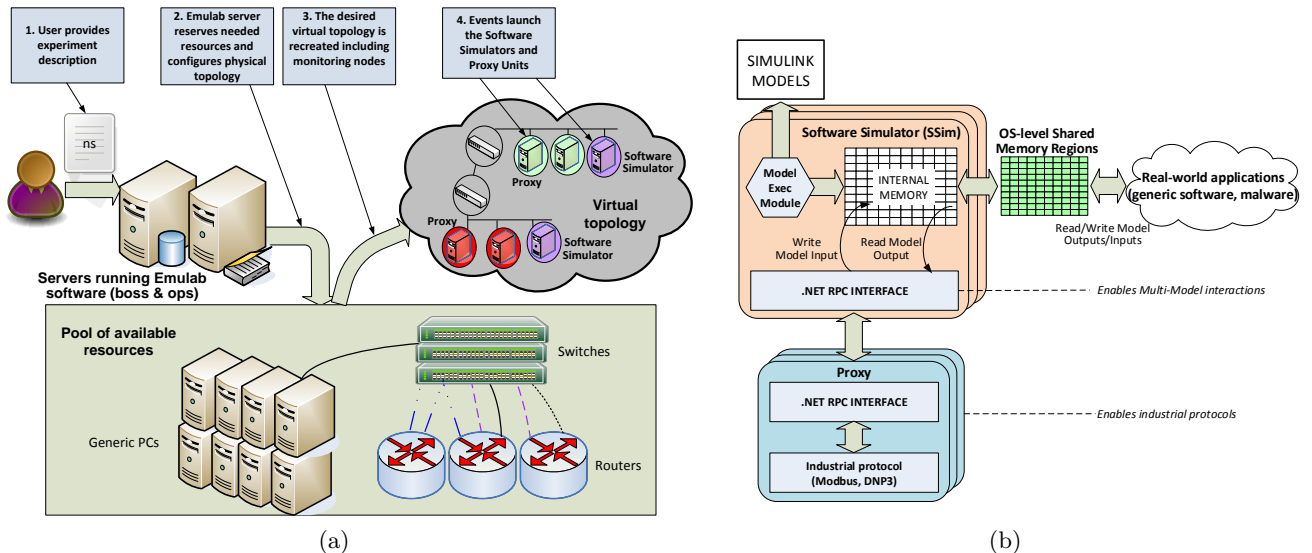


Figure 1: Architecture of the EPIC testbed: (a) Architectural overview and experimentation steps; and (b) EPIC software modules, including software simulators (SSim) and Proxy units.

ing Emulab in EPIC, we can automatically and dynamically map physical components, e.g., servers and switches, to a virtual topology. In other words, the Emulab software configures the physical topology in a way that it emulates the virtual topology as transparently as possible.

The basic Emulab architecture consists of two control servers, a pool of physical resources that are used as experimental nodes (generic PCs, routers or other devices) and a set of switches that interconnect the nodes. The Emulab software provides a Web interface to describe the steps that define the experiment life cycle within our testbed:

1. First we need to create a detailed description of the virtual network topology, *the experiment script*. The use of a formal language for experiment setup eases the recreation of a similar setting by any other researcher who wants to reproduce our results.
2. Experiments are instantiated by using the Emulab software, which automatically reserves and allocates the physical resources that are needed from the pool of available components.
3. The software configures network switches in order to recreate the virtual topology by connecting experimental nodes using multiple Virtual Local Area Networks (VLANs). Then, it configures packet capturing of pre-defined links for monitoring purposes.
4. Experiment-specific software, e.g., simulators, is launched automatically through events defined in the experiment script, or manually, by logging in to each node.

3.2 Recreating Physical Systems

Figure 1(b) provides an overview of the software units that recreate physical systems within EPIC. Physical process models are built in Matlab Simulink, from which the corresponding 'C' code is generated using Simulink Coder.

The generated code is then integrated in the software simulation unit (*SSim*) in order to enable real-time interaction of simulated processes with the rest of the emulation testbed.

At its core, the *SSim* unit provides the “glue” between the cyber and physical layers. From the *SSim*'s perspective models are “black-boxes” with inputs and outputs dynamically mapped to an internal memory region. Values written into this region are copied to the model's inputs, while model outputs are copied back to the internal memory. This way, EPIC enables experimentation with a wide range of physical processes without the need to provide any details on their actual content.

To enable interdependency studies on multiple NCIs, *SSim* implements a Remote Procedure Call (RPC) interface accessible to other *SSim* instances. RPC calls provide access to the internal memory region and consequently to the model's inputs and outputs, and enable real-time interactions between different models. Additionally, EPIC supports industrial protocols such as Modbus through *Proxy* units that translate calls between *SSim* and other units such as servers found in industrial installations.

3.3 Integrating Real-World Hardware and Software

Since almost everything is real, EPIC supports any software that usually runs on regular PCs and can practically integrate any hardware equipped with an Ethernet network interface. For instance, our installation includes real control hardware and real industrial software that enable studies on specific industrial architectures.

The interaction between real-world software and EPIC's software units is achieved in several ways. First of all, real software interacts with the simulated models using industrial protocols such as Modbus. Modbus calls are sent to a *Proxy* unit which forwards them as RPC calls to the *SSim* unit. Another way to interact with EPIC's software units is by Operating System (OS)-level shared memory. As shown in Figure 1(b), software units can access a shared memory

region that is mapped to the model’s inputs/outputs by the *SSim* unit. This technique enables interaction with software that does not implement RPC or Modbus and provides a simple way to run more complex security studies.

3.4 Real-time Simulation on Multitasking OS

Whenever real-time simulation is used, models run in a discrete time-domain that is closely linked to the clock of the OS. This means that the simulated model runs at the same rate as the actual physical system.

In EPIC we use generic PCs with multitasking OSs to run the real-time software simulation units. Our choice to use Simulink Coder to produce the simulators, although it has major advantages, imposes several constraints on the simulated models. An important aspect in this sense is the choice of the model execution rate, also known as the *simulation step*. The model’s internal dynamics limit the range of possible simulation steps. Before choosing the simulation step the researcher needs to verify that the output of real-time simulation reproduces as accurately as possible the real-world process. In parallel, the model execution time on a specific computer is limited by the model’s complexity and by the host’s processing power. As a general rule we can state that if the model’s execution time exceeds its simulation step, real-time simulation is not possible.

To test the limitations of software simulation in EPIC we experimented with several physical processes (see Figure 2). Here we mention small-scale processes such as Bell and Åström’s oil-fired 160MW electric power plant [1], which is based on the Sydsvenska Kraft AB plant in Malmö, Sweden, and the Tennessee-Eastman chemical process [9], which is also based on a real process, but the authors have introduced slight changes in order to protect the identity of reactants and products. The railway systems used throughout our experiments are based on the train models proposed by Rios and Ramos [19]. These take into account several aspects of real transportation systems such as weight, speed, acceleration, deceleration, and power consumption. Finally, we mention the IEEE suite of power grid case systems [22], which are currently used with EPIC. The 9-bus test case is the Western System Coordinating Council’s (WSCC) 3-machine 9-bus system, while the 30-bus, 39-bus and 118-bus test cases represent a portion of the American Electric Power System as of early 1960. These constitute realistic models which are well-established within the power systems community and provide a wide range of power system configurations.

The results in Figure 2 show that real-time software simulation is well suited for small and medium-scale models. However, software simulation is limited by the CPU speed and the size of the model. For instance, the IEEE 118 bus system is a complex model that includes 54 generators with frequency of 50Hz and a maximum simulation step of approximately 24ms. Since the model’s execution time on a 2.8GHz CPU is 155ms, real-time simulation is not possible in this case.

This is a well-known limitation of real-time software simulation, but it can be addressed in several ways. For instance, researchers might leverage parallel processing techniques such as GPU computing. Another approach is to use dedicated hardware simulators that are more powerful and specifically designed for simulations. However, these are still very expensive and in a multi-model environment they could

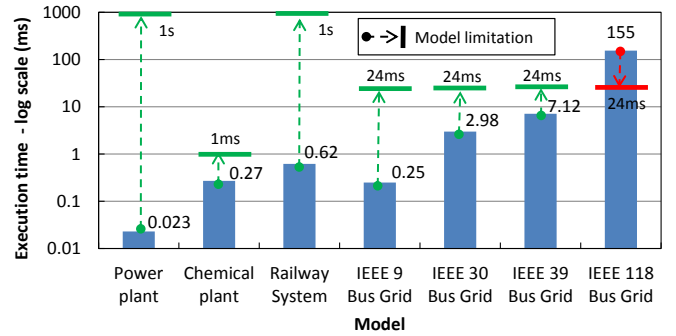


Figure 2: Execution time on a 2.8GHz CPU and limitations of various models. Here we see that EPIC enables experimentation with power plants, chemical plants, railway systems, and power grid models from the well-known suite of standard IEEE models. The red line on the IEEE 118 bus model highlights that real-time simulation requirements are not met since the model execution time exceeds the maximum simulation step allowed by the model dynamics.

render the cost of the cyber-physical testbed prohibitive.

3.5 Implementation Details

The installation of EPIC at the Joint Research Centre consists today of 120 PCs and approximately 100 virtual machines which are massively interconnected with two stacks of network switches. In addition carrier-grade routers, e.g., Cisco 6503, and industrial control hardware and software, e.g., ABB AC 800M control hardware including Modbus interfaces together with control server and Human Machine Interface software from ABB, are available as experimental resources. Software units such as *SSim* and *Proxy* have been developed in C# and have been ported and tested on Unix-based systems with the help of the *Mono* platform.

4. SCALABILITY AND APPLICABILITY TO REALISTIC SCENARIOS

EPIC’s ability to recreate both the cyber and physical dimensions of NCIs provides a large spectrum of experimentation options and enables complex experimentation with critical infrastructures. In this section we provide an overview of typical experiments conducted with EPIC and a description of a full experimental scenario.

4.1 Typical Experiments

Since 2009, EPIC has been concurrently used by many researchers for developing, testing, and validating a wide range of concepts, prototypes and tools. Table 2 provides a few examples of typical experiments that should give the reader a glimpse of the real applicability and scalability of EPIC. More details can be found in scientific reports and papers listed on the official Web site of EPIC (<http://ipsc.jrc.ec.europa.eu/?id=693>).

The first experiment is a study on the effect of network parameters to the effectiveness of cyber attacks targeting NCIs. More specifically, we evaluated the impact of network delays, packet losses, background traffic and network segmentation on a spoofing cyber attack consisting of an

Table 2: Examples of typical experiments performed with EPIC.

Objective	Nodes	Attack type	Physical process	Outcome
Evaluate the impact of network parameters (delay, packet loss, segmentation) on cyber attacks targeting NCIs	6–15	Spoofing	Power & Chemical Plants	The experiments proved that communications delays and packet losses have an insignificant effect on cyber attacks, but a physical process-aware network segmentation increases the resilience of NCIs. As a result, a novel segmentation approach was developed and validated.
Validate a novel approach for Anomaly Detection in cyber-physical systems	21	DDoS & Spoofing	Power Grid	The experiments lead to the validation of a proof-of-concept prototype showing the superior performances of combined cyber-physical Anomaly Detection Systems over traditional approaches that separate the cyber realm from the physical realm.
Evaluate the effect of operational decisions and their cascading effects on interdependent NCIs	12	Coordinated attack	Power Grid & Railway Transportation	The experiments illustrated the propagation of cyber disturbances (loss of monitoring & control) from one NCI to other dependent NCIs. They also proved that NCI operators need to collaborate in order to ensure the global stability of interconnected NCIs.
Analyze the effect of operator reactions and coordination in well-known (YouTube) Border Gateway Protocol (BGP)-route hijacking incidents	32	Hijacking & Man-in-the-middle	–	The study highlighted the importance of using adequate tools and mechanisms for a fast discovery of BGP hijacking events, the need for well trained operators with security expertise and the positive effect of a trusted medium that can support the communication and coordination between Network Service Providers (NSPs).

adversary capable to send legitimate commands to process controllers. The experiment was fully automated by EPIC to explore the parameter space and showed that while communications parameters have an insignificant effect on cyber attacks, a physical process-aware network segmentation can lead to more resilient systems.

The second experiment shows that studies can validate the effectiveness of newly proposed protection mechanisms. By using EPIC, we recreated a complex setting including real networks, protocols, hosts and routers, hence a realistic environment for the validation of a novel Anomaly Detection System (ADS). The experiment confirmed that EPIC can recreate a realistic environment in order to launch Distributed Denial of Service (DDoS) attacks together with spoofing attacks on critical infrastructure assets. These contributed to the validation of a novel ADS capable to efficiently detect anomalies both in the cyber and the physical dimension of NCIs.

The third experiment focuses on an important factor, i.e., the human operator, and closes a significant loop in cyber-physical experimentation. The experiment included a coordinated cyber attack in which the attacker prevented the normal remote operation of several substations, i.e., reduction of load, by blocking communications. Consequently, several substations exhibited a severe drop of voltages below nominal operating levels. The experiment demonstrated that operational decisions can make the difference between a complete breakdown and system survival and that collaborations between operators can limit the propagation of cyber disturbances.

Finally, the fourth experiment recreates the well-known YouTube Border Gateway Protocol (BGP)-route hijacking incident [18] and analyzes several hypothetical scenarios. During the experiment we developed an abstraction of Internet backbone networks in Europe and we recreated the infamous incident by replaying real traffic traces. The results highlighted the importance of adequate tools and mechanisms for a fast discovery of BGP hijacking events, and most importantly the need for well trained operators that can communicate over a trusted medium.

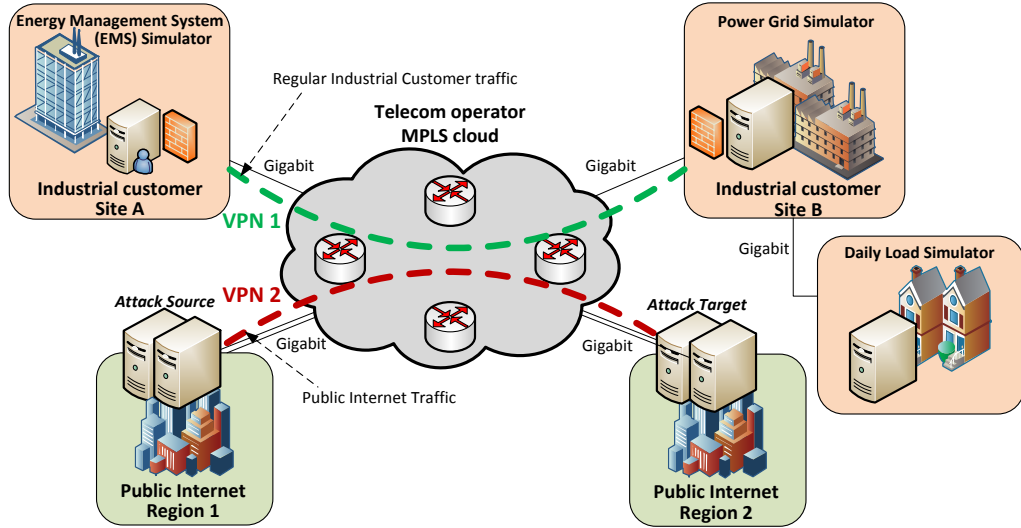
These experiments represent only a small fraction of the many directions and applications in which EPIC proved its value as a modern scientific instrument. The use of the platform is not constrained to disruptive experiments, but can be extended to educational and preparedness activities, e.g., as an environment for the execution of cyber security exercises.

4.2 Illustrative Experiment

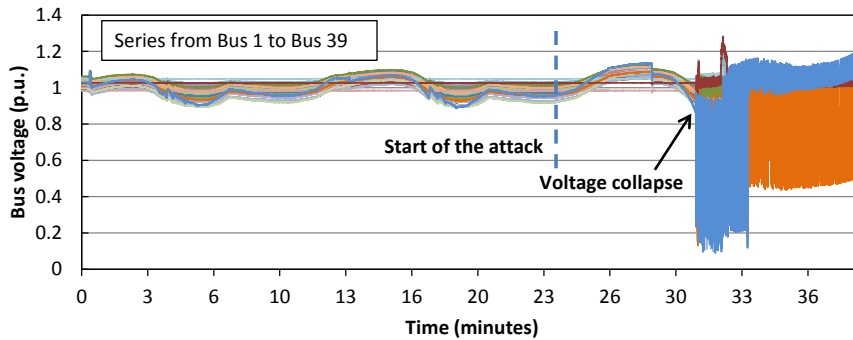
Next, we illustrate the applicability of EPIC by exploring what consequences ICT disruptions may have on the operation of a critical infrastructure such as the power grid. In this case study we consider the hypothetical scenario of a cyber-attack and specifically a DDoS attack, that is causing a severe telecommunication service degradation which propagates across critical infrastructures.

4.2.1 Experiment Setup

For the purposes of this experiment we recreated the typical architecture of an installation in which the power grid is



(a)



(b)

Figure 3: Impact of a cyber-attack on critical infrastructures: (a) Experimental setting with three physical system simulator *SSim* units, an Energy Management System simulator, attacker nodes and two virtual circuits offered by a telco operator (VPN 1: private circuit of the grid operator, VPN 2: public Internet); and (b) effect on voltage stability.

controlled remotely (Figure 3(a)). Here, *Site A* runs a simplified model of an Energy Management System (EMS) [21] to ensure voltage stability. The EMS continuously monitors and adjusts the operational parameters of the power grid model running at *Site B*. The commands are sent by the EMS to emulated control hardware, i.e., implemented through proxy units, that provide access to the power grid model inputs and outputs. Communications are provided by the Modbus protocol.

The IEEE electrical grid models are extensively used by the scientific community to conduct similar studies since they are known to accurately encapsulate the basic characteristics of real infrastructures. In our scenario we adopted the IEEE 39-bus New England system that includes a total of 39 substations together with 10 generators. The daily load imposed on our system derives from real data [13] and the intervention of the EMS is required to keep the grid stable.

To provide a realistic communications infrastructure between the EMS and power grid simulator we assumed that

the service provider uses an MPLS (Multi Protocol Label Switching) network. MPLS is a protocol that telco operators already use to replace older implementations based on Frame Relay and Asynchronous Transfer Mode (ATM) [12]. Using our Emulab installation, we created a minimal MPLS network with four Cisco 6503 routers, on which we defined two MPLS Virtual Private Networks (VPNs). VPN 1 acted as a protected virtual circuit between *Site A* and *Site B*, an approach that is usually followed by telco operators to isolate customer traffic. Since telco operators route diverse traffic, e.g., public Internet traffic, through the same MPLS cloud, we used VPN 2 to create a virtual circuit between two different “public” regions.

4.2.2 Telco Disruption and its Propagation to the Power Grid

Next, we launched a bandwidth consumption DDoS attack in VPN 2 and measured its effect on the power grid operator’s virtual circuit in VPN 1. The attack had serious

effects on the grid operator's private circuit. Consequently, the EMS lost control over the power grid and was unable to send commands that could restore stability. As shown in Figure 3(b), once the attack is started the grid is able to run for approximately 7 minutes without intervention. However, after 7 minutes the changes in the daily load would require the intervention of load shedding algorithms implemented within the EMS. Since the commands from the EMS can not reach the emulated control hardware, the voltages in the different segments of the grid begin to collapse.

Shortly after the attack is started the model becomes highly unstable and exhibits large oscillations which are difficult to map to reality. In fact, one of the major limitations of simulation-based studies is that we can only reason within the model's boundaries. However, voltage collapse is a clear indication of grid instability and in such cases operators might need to rebuild the entire grid. Therefore, we can state that for the purposes of our security study it suffices to verify that the attacker is able to lead the system outside the normal operating limits. If experiments need to go beyond these limits then researchers need to extend the models of physical systems to cover extreme and unstable conditions or to extend the cyber-physical testbed with real physical devices, if this is feasible, economically possible and safe.

4.2.3 A Look at Reality

In reality, most telco operators take strong measures to limit the interference between separate VPNs. For example, with the deployment of Quality of Service (QoS) in the MPLS network an attack on the public Internet hardly affects the private traffic of other telco service customers. This claim was actually validated by running the aforementioned experiment after activating QoS (with packet prioritization - a feature that might also be used to implement packet prioritization in industrial communications) in the MPLS cloud. The only measurable effect was a slight increase of the packets Round-Trip-Time (by 1-2ms) - a tolerable delay if we consider the IEEE 1646-2004 standard for communication delays in substation automation, which states that high-speed messages must be delivered in the 2ms to 10ms range.

Nevertheless, such protective measures are not compulsory, e.g., through policies and regulation. The severe risks that are involved if such protective measures are not implemented, were clearly demonstrated by our case study that highlights the potential impact of ICT disruptions on a wide range of physical systems.

Furthermore, by designing and conducting experiments using evidence from real incidents we can effectively explore several "what-if" scenarios. As an example, we consider an incident that occurred on January 2, 2004 and was related to Rome's remotely controlled power grid [4]. During this incident communications between remote sites were disrupted due to a broken water pipe that flooded the server room of a Telecom operator and short-circuited critical hardware. Consequently, power grid operators were completely blinded and could not monitor or control the remote site. Fortunately, there were no disturbances, so the grid remained stable. Nevertheless, as shown by experiments we performed on EPIC, a change in the balance between generated and consumed energy, would have serious consequences on the electrical grid. In Rome, this could have led to black-outs

throughout the entire city, and affect other critical infrastructures as well like transportation and health-care systems.

5. CONCLUSIONS

By combining an Emulab-based testbed with real-time software simulators, EPIC demonstrates a novel approach to conduct cyber security studies with multiple heterogeneous NCIs. EPIC can be considered as an instance of a new class of scientific instruments, namely cyber-physical testbeds, that are suitable for assessing the impact of cyber-threats against physical infrastructures. Our experience with EPIC proves that such instruments can support interesting studies in many interdependent critical infrastructure sectors and with heterogeneous systems such as transportation networks, power plants, chemical plants, and power grids. Few of these studies have been recorded as Demo videos and are available at the following link: <http://ipsc.jrc.ec.europa.eu/?id=691>.

References

- [1] Bell, R. and Åström, K. Dynamic models for boiler-turbine alternator units: data logs and parameter estimation for a 160MW unit. *Lund Institute of Technology, Report TFRT-3192*, 1987.
- [2] Benzel, T., Braden, R., Kim, D., Neuman, C., Joseph, A., Sklower, K., Ostrenga, R., and Schwab, S. Experience with DETER: A testbed for security research. In *Proc. of the Int. Conf. on Testbeds and Research Infrastructures for the Development of Networks & Communities (TridentCom)*, 2006.
- [3] Bergman, D. C., Jin, D., Nicol, D. M., and Yardley, T. The virtual power system testbed and inter-testbed integration. In *Proceedings of the 2nd conference on Cyber security experimentation and test, CSET'09*, pages 5-5, Berkeley, CA, USA, 2009. USENIX Association. URL <http://dl.acm.org/citation.cfm?id=1855481.1855486>.
- [4] Bobbio, A., Bonanni, G., Ciancamerla, E., Clemente, R., Iacomini, A., Minichino, M., Scarlatti, A., Terruggia, R., and Zendri, E. Unavailability of critical SCADA communication links interconnecting a power grid and a telco network. *Reliability Engineering & System Safety*, 95(12):1345-1357, 2010.
- [5] Charette, R. IT Hiccups of the Week: Southwest Airlines Computer Failure Grounded All Flights. *IEEE Spectrum*, June 2013.
- [6] Chen, T. and Abu-Nimeh, S. Lessons from Stuxnet. *Computer*, 44(4):91-93, april 2011.
- [7] Chertov, R., Fahmy, S., and Shroff, N. B. Fidelity of network simulation and emulation: A case study of TCP-targeted denial of service attacks. *ACM Trans. Model. Comput. Simul.*, 19(1):4:1-4:29, 2009.
- [8] Chunlei, W., Lan, F., and Yiqi, D. A simulation environment for SCADA security analysis and assessment. In *Proc. of the 2010 International Conference on Measuring Technology and Mechatronics Automation*, pages 342-347, 2010.

- [9] Downs, J. and Vogel, E. A plant-wide industrial process control problem. *Computers & Chemical Engineering*, 17(3):245–255, 1993.
- [10] Duggan, D. Penetration testing of industrial control systems. *Technical Report, SAND2005-2846P, Sandia National Laboratories*, 2005.
- [11] Hahn, A., Ashok, A., Sridhar, S., and Govindarasu, M. Cyber-physical security testbeds: Architecture, application, and evaluation for smart grid. *Smart Grid, IEEE Transactions on*, 4(2):847–855, 2013.
- [12] IBM and Cisco. Cisco and IBM provide high-voltage grid operator with increased reliability and manageability of its telecommunication infrastructure. *IBM Case Studies*, 2007.
- [13] Manera, M. and Marzullo, A. Modelling the load curve of aggregate electricity consumption using principal components. *Environ. Model. Softw.*, 20(11):1389–1400, November 2005.
- [14] M.J. McDonald and J. Mulder and B.T. Richardson and R.H. Cassidy and A. Chavez and N.D. Pattengale and G.M. Pollock and J.M. Urrea and M.D. Schwartz and W.D. Atkins and R.D. Halbgewachs. Modeling and simulation for CyberPhysical system security research, development and applications. *Technical Report, SAND2010-0568, Sandia National Laboratories*, 2010.
- [15] Nai Fovino, I., Masera, M., Guidi, L., and Carpi, G. An experimental platform for assessing SCADA vulnerabilities and countermeasures in power plants. In *Human System Interactions (HSI), 2010 3rd Conference on*, pages 679–686, 2010.
- [16] Nan, C., Eusgeld, I., and Kröger, W. Analyzing vulnerabilities between SCADA system and SUC due to interdependencies. *Reliability Engineering & System Safety*, 113(0):76–93, 2013.
- [17] Queiroz, C., Mahmood, A., and Tari, Z. SCADASim – a framework for building scada simulations. *Smart Grid, IEEE Transactions on*, 2(4):589–597, 2011.
- [18] R. NCC. YouTube Hijacking: A RIPE NCC RIS case study. <http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>, 2008. [Online; accessed April 2013].
- [19] Ríos, M. A. and Ramos, G. Power system modelling for urban massive transportation systems. *Infrastructure Design, Signalling and Security in Railway*, pages 179–202, 2012.
- [20] Siaterlis, C., Garcia, A., and Genge, B. On the use of Emulab testbeds for scientifically rigorous experiments. *IEEE Communications Surveys and Tutorials*, 15(2): 929–942, 2013.
- [21] Tuan, T., Fandino, J., Hadjsaid, N., Sabonnadiere, J., and Vu, H. Emergency load shedding to avoid risks of voltage instability using indicators. *Power Systems, IEEE Transactions on*, 9(1):341–351, feb 1994.
- [22] University of Washington - Electrical Engineering. Power Systems Test Case Archive. <http://www.ee.washington.edu/research/pstca/>, 2012. [Online; accessed April 2013].
- [23] US Department of Energy. National SCADA Test Bed. http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/NSTB_Fact_Sheet_FINAL_09-16-09.pdf, 2009. [Online; accessed April 2013].
- [24] White, B., Lepreau, J., Stoller, L., Ricci, R., Guruprasad, S., Newbold, M., Hibler, M., Barb, C., and Joglekar, A. An integrated experimental environment for distributed systems and networks. In *Proc. of the 5th Symposium on Operating Systems Design and Implementation*, pages 255–270, 2002.
- [25] Yardley, T., Berthier, R., Nicol, D., and Sanders, W. Smart grid protocol testing through cyber-physical testbeds. In *Proc. of the 4th IEEE PES Innovative Smart Grid Technologies (ISGT 2013) Conference*, 2013.

Christos Siaterlis (christos.siaterlis@jrc.ec.europa.eu) is a project officer at the Institute for the Protection and Security of the Citizen of the European Commission’s Joint Research Centre, Ispra, Italy.

Béla Genge (bela.genge@ing.upm.ro) is a Marie Curie post-doctoral fellow and a member of Informatics department at Petru Maior University of Tg. Mureş, Mureş, Romania.