

A VIRTUAL POWER SYSTEM TESTBED FOR CYBER-SECURITY DECISION SUPPORT

David M. Nicol

Department of Electrical & Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL, 61801, U.S.A.

Charles M. Davis

PowerWorld Corporation
Champaign, IL 61801, U.S.A.

Thomas Overbye

Department of Electrical & Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL, 61801, U.S.A.

ABSTRACT

This paper describes a testbed for evaluating power grid cyber-security and its support for decision-making. The testbed combines physical hardware and software components used in the grid with an electric generation/distribution simulator, and a computer/communications simulator. We describe the relationship between these three parts, and two scenarios where the testbed might be used to support decision making.

1 INTRODUCTION

The electric power grid is controlled by an intricate computer and communication infrastructure. Cyber-security of this infrastructure is an issue of paramount importance, with many pressing questions : what sort of security architecture best balances competing concerns of availability, safety, and security? What impact does any given security technology have on the operators, and on the system's ability to keep up with real-time monitoring requirements? How does one balance the cost of implementing and maintaining security measures against the risk of not doing so? Answers to questions like these are essential for decision makers faced with the task of securing their portion of the power grid.

The Trustworthy Cyber-Infrastructure for Power (TCIP) project at the University of Illinois (with partners at Cornell University, Dartmouth College, and Washington State University) is funded by the National Science Foundation, Department of Energy, and Department of Homeland Security to study cyber-security issues in the electric power

grid (UIUC 2008). A central component of the project is the Virtual Power System Testbed (VPST), used to model and simulate large portions of the grid. Through generous donations by industrial sponsors VPST has real hardware and real software used in the monitoring and control of the grid. The real devices are seamlessly integrated with a simulator that adds virtual devices and networking to the model; the real devices and computer/communications simulator are both integrated with a simulation of electric power generation and distribution. In normal operation, real and virtual sensing devices are queried by real and virtual control functions for electrical system state, using standardized protocols that are deployed in the grid. Results of these queries are returned to real control station software which places the results in a database and displays them to an operator.

The real power of VPST lies in the scale of the systems it can model, and in the ways it can be used to study properties of modeled systems. All of the communication between devices (real and virtual) that occurs in the modeled system is carried by the communication system simulator. Cyber-attacks and cyber-defenses can be simulated within the communication system simulator, with the attendant impact on the power grid monitoring and control traffic, and attendant impact on the operators' situational awareness. The ability of a system to defend against and even tolerate intrusions can be evaluated within the context of VPST. Correspondingly, VPST can be used in a decision support role both at design time (e.g., what sort of architecture ought we to deploy) and operationally (e.g., on suspecting a cyber-

penetration, what impact may various response options have on availability?).

We turn first to a description of the three components of VPST and how they fit together, and then sketch use case scenarios supportive of decision-making.

2 VSPT Components

2.1 PowerWorld

The PowerWorld simulator is an outgrowth of research done at the University of Illinois, and has been commercialized. It is used in hundreds of power control centers around the world to predict the equilibrium state of the electric grid as a function of the load (i.e., demand) and generation. PowerWorld represents generation stations, electrical lines (and their capacities), and “buses”—points in the grid where power lines come together, with some sort of transformation applied (e.g., a voltage transformers). In the physical infrastructure buses occur within *substations*, where sensors are placed to monitor the electrical state, where computers gather that state and report it on demand to a control station. PowerWorld is capable of simulating up to 100,000 buses—a city the size of Madison, Wisconsin contains a couple dozen substations and a couple hundred buses. PowerWorld runs on an ordinary PC, and gives its users a graphical interface to the solution engine, including animation of the power flows and the state of the buses in the model. In our discussion we will use VPST-E to refer to the version of PowerWorld that was modified for VPST.

2.2 Real Components

The VSPT has several *relays*. In normal operation these sense the current on a bus, and open a *breaker* to interrupt the electrical flow at that point if a fault condition (e.g. a short circuit) is detected. Relays (and other sensors, all of which are generically called *remote terminal units (RTU)*) sense electrical state data, and communicate that on demand to a *data aggregator*. In one (of many) typical substation architectures there is one data aggregator, serving several RTUs. In the past substations used dedicated serial communication lines (e.g. RS232) to connect the data aggregator and RTUs, new IP-based technology is increasingly used, both wire-line and wireless. The devices use protocols such as *Modbus* ([Modbus.org 2005](http://Modbus.org)) and *DNP3* ([DNP.org 2008](http://DNP.org)) to communicate. In such a connection one device acts as the master, which sends query commands to the slaves. In a substation the data aggregator is master, the RTUs are slaves, and the query typically is asking for the present set of sensor values observed at the RTUs.

The VSPT also has *control stations*, and *data historians*. A control station is software that monitors substations, and provides an interface to its operator to send control signals to

the power grid devices. A data historian is a database where the queried state information is placed for examination by the substation (and other systems in the enterprise. Here the control station (through a device called a *front-end processor*) uses Modbus or DNP3 (or some other standardized protocol), acting as the master, and queries substations’ data aggregator (acting in this transaction as a slave) for electrical state. Historically the communication technology used for this interaction is over leased telephone lines. These too are rapidly being replaced by ethernet based communication.

A VPST model may have other real devices that are not part of the power control system, e.g., hosts that run cyber-attacks on the simulated system, or a web-server that VPST devices query. We use the label VPST-R to refer to the collection of all real devices and software stacks that are part of a VPST model.

2.3 Computer/Communication Simulation

The VPST uses a very capable detailed simulator of computer hardware, software, and communication infrastructure, called here VPST-C. Based on RINSE ([Liljenstam, Nicol, Yuan, Yan, and Liu 2006](#)) and PRIME ([Liu 2008](#)), VPST-C is notable in its ability to execute models of traffic and architectures at multiple levels of resolution, concurrently; it is likewise unique in its ability to integrate both purely virtual simulation models, and emulation models where real traffic is sent or received by the same software that runs on real devices, but the traffic is carried by the network simulation. VPST-C can simulate a variety of wireline and wireless technologies used in the power grid. The simulator represents every device in the model as a software stack modeled after the ISO network stack ([Kurose and Ross 2007](#)); a purely virtual device has simulation modules from the application layer all the way down to the physical layer, while an emulated host has a proxy representation inside the simulator with an “emulation layer” at the top, followed by the same modules for layers 3,2, and 1 that the purely virtual hosts use. Traffic from the emulated host is captured and sent to VPST-C. The source’s real IP address is transformed to the one used by that device’s proxy inside of the simulator, and is presented to the emulation layer. From there the traffic is pushed into the simulated communication infrastructure. The opposite occurs when traffic is delivered to an emulated host’s proxy : the traffic rises through layers 1,2, and 3 of the proxy, is delivered to the emulation layer, and is pushed back out to the real software stack.

VPST-C has models of a variety of cyber-attacks (e.g., distributed denial of service, worms), and models of technologies used in cyber-defensive (e.g. bump-in-the-wire cryptography, detection of back-scatter from worm scans).

3 VPST Component Integration

Each of the three VPST components share the same architectural view of the system being modeled; a common textual name space exists for referring to the devices. Each system has its own internal naming mechanisms, the common name space simply supports translation between different systems. Of course, some architectural details of importance to one simulator have no bearing on the other and need no translation. A high level depiction of the three components and

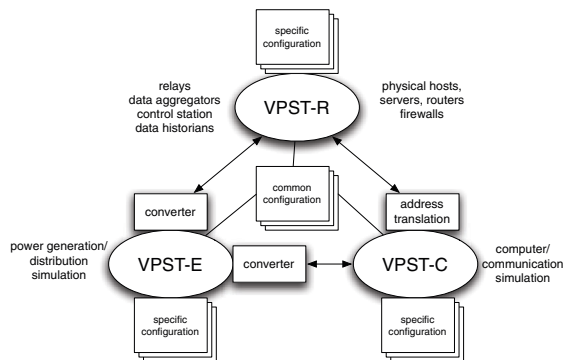


Figure 1: Components of the Virtual Power System Testbed.

Connection between VPST-E and VPST-R is limited. The heart of the integration is to provide relays within the VPST-R with descriptions of electric flows at the points in the VPST-E topology where the relays are positioned in the virtual system. Real relays sense and react to electrical *current*; because field currents are not seen in laboratories and testbeds, relays can also be configured to sense and react to *voltage* that emulates (in an analog sense) the current. This functionality is normally provided for testing purposes, but is something we exploit in VPST. A programmable device known as an adaptive multi-channel source (AMS) is used to translate a digital description of a three phase flow into sinusoidal voltage waves connected to the relay. The relay makes its control decisions based on these voltages, just as it would in the field, based on the currents modeled by the voltages. A converter program (normally on the same PC as PowerWorld) uses TCP-based sockets to get from PowerWorld simulated current values. The converter then programs the AMS to output corresponding 3 phase voltage to the relay. The connection between converter and AMS is a serial line.

A much more extensive convertor process exists between the VPST-E and VPST-C : every simulated RTU within the VPST-C needs the values the VPST-E produces at its point in the grid topology. The common name space for buses allows VPST-C to take a list of buses and their electrical state and map it to where the simulated RTUs can reference it when queried for electrical state.

Finally, the connection between routable devices within VPST-R and VPST-C is handled through normal routing and address translation mechanisms applied to IP streams. VPST-C contains an OpenVPN server configured to obtain packets from VPN clients installed on VPST-R devices; those devices that do not run the VPN client explicitly redirect their traffic to a simple proxy server within VPST-C. In either case packet flows involving devices in the VPST-R have their traffic pass through VPST-C servers that remap those flows to VPST-C.

4 Use Cases

VPST has supports decision-making in a variety of ways, illustrated by simple use cases.

4.1 Decision Recognizing Need

The first and most basic decision is to recognize the need to secure the power grid cyber-infrastructure. VPST has been used several time to demonstrate vulnerabilities, as we next outline.

The power grid is a real-time control system; control requires the operators to have up-to-date *situational awareness*, i.e., up-to-date knowledge of what the electrical system is doing. Lack of situational awareness was identified as a root cause of the massive mid-western blackout in August 2003 (PNNL 2003). One way to deny situational awareness is to mount a cyber-attack on the communication infrastructure. Power control networks are integrated into the larger networked enterprise, and this offers opportunities for intrusion from outside the enterprise. One use case scenario we demonstrated on VPST used the following steps:

1. A substation worker using a vulnerable PC visits a web-site that infects the PC, called now the “zombie-master”.
2. The zombie-master reaches back to the Internet and downloads a reconnaissance package.
3. The recon software maps the control network, identifies live IP addresses, and other vulnerable PCs, and creates zombies of the other vulnerable PCs with malware that participates in a bandwidth consumption attack, on command.
4. The zombie-master reports its identity and results to the attacker, and awaits command to launch the distributed denial of service attack.
5. The attacker changes the flow topology of the electric grid by logging into relays identified by the initial scan, using default passwords that have been left in place.
6. The attacker simultaneously commands the zombie-master to launch the bandwidth consumption attack.

7. As a result of the DDoS attack, the communication between control station and substations is interrupted, and the operators lose all situational awareness. They know that they cannot see the electrical state, but have no means of observing it.

In this scenario the control station, one relay and one data aggregator were used in the VPST-R. Real DNP3 traffic was generated and responded to. All the details of the cyber-attack were completely simulated, although any one of them could have been emulated if we had used real scanning and real attack software.

We have also used the VPST to demonstrate a more subtle form of denying situational awareness. The Modbus protocol is very simple, and carries no sequence numbers or time-stamps. If a Modbus master receives a response to a query within a short period of time, it assumes that response is from the device it queried, with state information that is current at least to the time when the query was sent. However, well-known attack techniques can place a “man-in-the-middle” between the master and slave, so that the rogue device in the middle responds to the master’s queries with whatever state information it might choose to provide. This allows the rogue device to report a completely fabricated view of the electrical state, or simply a *delayed* view of the state. In the latter form of attack the rogue can forward master queries to a slave as though the rogue were the master, and store the responses. The rogue can build up a temporal cache of electrical state information, and when the real master queries for it, the rogue can return *old* state information. This helps the returned data survive consistency checks that the real master might apply to the data, but still keep the operator from having a true picture of the state of the system.

4.2 Decision to adopt or not adopt technologies

Introduction of security technologies can impact the length of messages (which now may carry extra bits to support authentication), the path of messages (e.g., to pass through a filtering device), and the number of messages (again to support authentication protocols). The extra overhead exacts a toll on a message’s latency, and on the supportable bandwidth of control and monitoring messages. We have used VPST in contexts where the performance cost of a considered technology is assessed to determine whether the additional overhead brings the latency too close to real-time deadlines. One technology we evaluated called for the RTUs to run authentication protocols on all messages, and pass them through a centralized security hub, which would then redistribute them. The design was intended to thwart any unregistered device from speaking on the network, and to simplify the problems of cryptographic key management. The study suggested the approach was viable

for sub-stations with a small-to-moderate number of RTUs, but suffered from scaling problems for the larger substation designs.

Another study uses the VPST to assess the impact on message latency of using so-called “bump-in-the-wire” (BiW) cryptography devices. These may be placed on legacy serial communication lines at both the point of transmission and point of receipt. A cryptographic hash is computed for a message on its transmission, and checked upon receipt. If the message is intercepted and altered, or if a rogue message is somehow inserted onto the channel, the recipient BiW device detects this and flags the message as having failed a CRC redundancy check. The message’s recipient sees the failure, and responds as though the message simply suffered a transmission error.

4.3 Decision on response mechanism

The VPST can simulate a variety of cyber-attacks, including real ones that are mounted from devices in the VPST-R on the rest of the system. The VPST can run intrusion detection systems that create alerts when malicious activity is detected. For example, an intrusion detection system could passively monitor Modbus traffic, and look for inconsistencies in the series of Modbus queries/responses, e.g., a query from a previously unknown device now posing as a master could trigger an alert (and might do so in the man-in-the-middle attack described earlier). There are a variety of responses an operator might take to such an alert, ranging from doing nothing, to implementing some sort of active filter on the traffic, to isolating the control network entirely. Each decision has different ramifications on the operation and performance of the system. A potentially fruitful role for VPST would be as a tool to investigate the utility and costs of different responses, and so be able to analyze the alternatives and choose the one that best maximizes the decision-maker’s utility. We envision VPST being used in this way first as a planning aid, but see the potential for using this kind of technology to support real-time decision-making to respond to suspected cyber-attacks.

5 Summary

It is crucial that cyber-security be brought to the cyber-infrastructure that controls the electric power grid. There are many decisions to be made concerning whether to protect, what to protect, how much to protect, what to do when some intrusion is suspected. The Virtual Power System Testbed is a unique tool that enable decision makers to observe the consequences of potential decisions, in a safe virtual environment. VPST achieves its flexibility and power through a unique integration of real devices, simulation of electric power generation and distribution, and simulation/emulation of computer and communication de-

vices. This paper describes these components, how they are integrated, and several use case scenarios that illustrate how VPST supports decision making. We anticipate integrating VPST further with other testbeds, and ultimately becoming a resource for others to use in the study of cyber-security in the power grid.

ACKNOWLEDGMENTS

This material is based in part upon work supported by National Science Foundation under Grant No. CNS-0524695. 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 National Science Foundation.

REFERENCES

- DNP.org 2008. DNP: Distributed network protocol. <http://www.dnp.org>.
- Kurose, J., and K. Ross. 2007. *Computer networking: A top down approach*. Addison-Wesley, 4th Edition.
- Liljenstam, M., D. Nicol, Y. Yuan, G. Yan, and J. Liu. 2006, Jan.. Rinse: the real-time interactive network simulation environment for network security exercises. *Simulation : Transactions of the Society for Modeling and Simulation International* 82 (1): 43–59.
- Liu, J. 2008, April. A primer for real-time simulation of large-scale networks. In *Proceedings of 41st Annual Simulation Symposium*, 85–94. Ottawa, CA.
- Modbus.org 2005. Modbus ida.
- PNNL 2003. Looking back at the august 2003 blackout. <http://eioc.pnl.gov/research/2003blackout.stm>.
- UIUC 2008. Trustworthy cyber-infrastructure for the power grid. <http://www.iti.uiuc.edu/tcip>.

AUTHOR BIOGRAPHIES

DAVID M. NICOL received a B.A. in math from Carleton College (1979), and a Ph.D. in computer science from the University of Virginia (1985). He is currently Professor of Computer and Electrical Engineering at the University of Illinois at Urbana-Champaign. His research interests include high performance computing, modeling and discrete-event simulation, and security. He is a Fellow of the IEEE, and Fellow of the ACM. His e-mail address is [<dmnicol@illinois.edu>](mailto:dmnicol@illinois.edu)

CHARLES M. DAVIS received the B.S. degree in electrical engineering from Louisiana Tech University in 2002, the M.S. degree (2005) and Ph.D. (2009) from the University of Illinois Urbana-Champaign. Currently he is working for the PowerWorld corporation. His current research interests include linear sensitivities, power system analysis, power

system visualization, and power system operational reliability. His e-mail address is [<matt@powerworld.com>](mailto:matt@powerworld.com).

TOM OVERBYE received the B.S., M.S. and Ph.D. degrees in electrical engineering from the University of Wisconsin-Madison. He is currently the Fox Family Professor of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. He was with Madison Gas and Electric Company, Madison, WI, from 1983-1991. His current research interests include power system visualization, power system analysis, and computer applications in power systems. He is a Fellow of the IEEE. His e-mail address is [<overbye@illinois.edu>](mailto:overbye@illinois.edu).