SANDIA REPORT

SAND2008-0910
Unlimited Release
January 2008

Towards Risk-Based Management of Critical Infrastructures: Enabling Insights and Analysis Methodologies from a Focused Study of the Bulk Power Grid

Randall A. LaViolette, Bryan T. Richardson, and Benjamin K. Cook

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>

Towards Risk-Based Management of Critical Infrastructures: Enabling Insights and Analysis Methodologies from a Focused Study of the Bulk Power Grid

Randall A. LaViolette¹, Bryan T. Richardson², and Benjamin K. Cook¹

¹Advanced Analytic Concepts Development, ²Critical Infrastructure Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-1235

Abstract

This report summarizes research on a holistic analysis framework to assess and manage risks in complex infrastructures, with a specific focus on the bulk electric power grid (grid). A comprehensive model of the grid is described that can approximate the coupled dynamics of its physical, control, and market components. New realism is achieved in a power simulator extended to include relevant control features such as relays. The simulator was applied to understand failure mechanisms in the grid. Results suggest that the implementation of simple controls might significantly alter the distribution of cascade failures in power systems. The absence of cascade failures in our results raises questions about the underlying failure mechanisms responsible for widespread outages, and specifically whether these outages are due to a system effect or large-scale component degradation. Finally, a new agent-based market model for bilateral trades in the short-term bulk power market is presented and compared against industry observations.

ACKNOWLEDGMENTS

A number of Sandia staff and external collaborators contributed significantly to the research documented in this report. In particular, the authors are grateful for the assistance given by Tu-Thach Quach, Jason E. Stamp, Walt A. Beyeler, Lori A. Ellebracht (TVA), Kevin L. Stamber, Edward V. Thomas, Satish Ranade (NMSU), Christopher Duran (formerly NMSU), David G. Robinson, George E. Apostolakis (MIT), Anthony M. Koonce (formerly MIT), C. J. Giesler (RESPEC), Michael G. Stickland, Robert J. Glass, David J. Borns, Laurence R. Phillips, and Hamilton Link. This work was funded by Sandia's Laboratory Directed Research and Development Program (LDRD) through the Energy and Homeland Security Investment Areas.

CONTENTS

1.	Introduction	7
1.1	The Problem	7
1.2	Research Approach and Objectives	9
2.	Physical Grid Model Development	11
2.1	Background	11
2.2	Theory	11
2.3	Software Architecture and Implementation	12
2.4	Verification and Validation	15
2.5	Usage	15
3.	Electric transmission grid vulnerability to cascade failures	17
3.1	The AC Electric Power (ACEP) Model	19
3.2	Results	22
3.3	Discussion and Conclusions	25
4.	Price Divergence in a Model of Short-Term Electric Power Market with Unsupervised Bilateral Trades between Myopic Agents	27
4.1	Introduction	27
4.2	Background	28
4.3	Description of the Model	29
4.4	Results	33
4.5	Discussion and Conclusions	38
4.6	Appendix A	39
4.7	Appendix B	41
4.8	References	42
5.	Conclusions	43
6.	Distribution	45

FIGURES

Figure 1.	Conceptual model of the bulk power grid	11
Figure 2.	Interpretations of power outage distribution	18
Figure 3.	Another interpretation of power outage distribution	19
Figure 4.	WECC transmission grid from anonymized 2005 data	21
Figure 5.	Distribution of load lost from model results	24
Figure 6.	Comparison of Hill statistic for an exponential and a Pareto distribution	25
Figure 7.	Hill statistic measured for the worst case load lost	25
Figure 8.	Schematic of contracted price to deliver power	27
Figure 9.	Supply-demand curves of the agent's initial conditions for a typical trading session	31
Figure 10.	Log-log plot of the sales price evolution generated in a typical trading session	34
Figure 11.	Sales price evolution of a typical trading session	35
Figure 12.	Maximum price evolution of the trading sessions	36

Figure 13. Comparison of model results to industry sales price data	37
Figure 14. Log-log plot of the empirical density distributions of the sales price	37

TABLES

Table 1. Statistics for load lost from model results..	23
--	----

1. INTRODUCTION

Our way of life in the U.S. depends upon a complex interdependent system of infrastructures. These infrastructures are currently vulnerable to disruptions that can lead to cascading failures with serious consequences. Our future security and wellbeing depend upon the successful application of science and technology to develop robust infrastructures that can withstand both terrorist and natural threats. As an integrated human-cyber-physical system, a resilient “living” infrastructure of the future could perceive perturbations, analyze the potential impact of changes to its system health, and respond in a timely manner to minimize the likelihood of realizing some undesirable state or consequence.

The purpose of this LDRD project was to develop analysis tools and enabling insights that would allow us to someday realize this vision of living infrastructures. The premise for our research was that effective infrastructure management, be it real-time active control using a network of distributed sensors and software agents or long-term evolutionary approaches based on policy and regulation, requires a systematic approach to identify the significant relevant risks from potential infrastructure failures and evaluate remedial alternatives. Current risk assessment techniques, like the N-1 contingency studies performed by electric utilities to assess the reliability of their systems, are unable to capture the growing complexity of our critical infrastructure systems that arises from their increasing scale, interconnectivity, and utilization.

To address this challenge, this research formulated a representative and holistic analysis framework to assess and manage risks in complex infrastructures. Our work was grounded through the detailed study of a specific infrastructure: the bulk power grid. We choose to focus on the power grid because it is arguably the most complex critical infrastructure due to its sheer size and highly dynamic nature. A comprehensive computer model of the bulk power grid was conceptualized that would approximate the coupled dynamics of its physical, control, and market components. This model was partially implemented and then integrated into a new risk management methodology to support improved decision making of mitigation options based upon the overall utility of decisions from the perspective of a decision maker. Significant progress was made in this project towards the development of the comprehensive grid model with the implementation and application of a novel power flow simulator that includes control elements such as relays. An agent-based market model for bilateral trades in the short-term bulk power market was also formulated and tested. This paper documents the power grid and market models and their application results; companion work on an overarching bulk power grid risk analysis methodology is described elsewhere¹.

1.1 The Problem

Our critical infrastructure is characterized by a complex, diffuse, and interdependent “system-of-systems” that is dominated by poorly understood complexities. Complexity in infrastructures

¹A.M. Koonce, G.E. Apostolakis, and B.K. Cook, Bulk power grid risk analysis: ranking infrastructure elements according to their risk significance, MIT Engineering Systems Division Working Paper ESD-WP-2006-19, 2006 (<http://esd.mit.edu/wps/esd-wp-2006-19.pdf>). See also final version of this paper to appear in *International Journal of Electrical Power and Energy Systems*.

systems is introduced at varying scales through the hierarchy of subsystems comprising physical and cyber components, linkage of these components within and across infrastructures, as well as human and economic interfaces. Driven by increased economic efficiencies, infrastructure components and systems have become increasingly more coupled and controlled through a cyber infrastructure layer often implemented on the relatively insecure Internet. Infrastructures are hybrids of old and new subsystems, which have myopically evolved largely with only local operational concern for either security or system vulnerabilities.

The electric infrastructure in the U.S. offers a good example of this complexity. Having started as isolated individual power systems supplying electricity to local regions across the country, the electric infrastructure evolved over the 20th Century to become a highly interconnected and interdependent system spanning much of North America. The electric infrastructure consists of three parts: the generation of electric power, the transmission of electricity, and the distribution of electricity to the end-users. The generation and transmission components are referred to as the bulk power system, which is the focus on this study and can be thought of as the backbone of the electric infrastructure. The bulk power grid is an international system that is divided into three major regions. These regions are collectively known as the NERC (North American Electric Reliability Council) Interconnections and include the Eastern Interconnection, the Western Interconnection, and the ERCOT (Electric Reliability Council of Texas) Interconnection. The Eastern Interconnection supplies power to the U.S. states and Canadian provinces east of and including the Great Plains region. The Western Interconnection supplies power to states and provinces west of and including the Rocky Mountain area. The smallest interconnection, ERCOT Interconnection, covers the majority of Texas. The infrastructure representing these interconnections is strongly connected within the interconnection regions but only weakly connected across the regions. Over the past several decades the grid has become increasingly automated with an array of networked, computer-controlled sensors and actuators² managing various aspects of the generation and flow of power through the physical infrastructure. As will be explained in more detail in Section 4 of this report, the physical interconnection of utility systems across the country combined with deregulation of the power industry in the 1970's and 1990's has fostered the rise of a complex bulk power market, through which wholesale power is bought and sold between utilities.

Our economy, national security, and very way of life are highly dependent on the availability and reliable supply of electric power. As the grid has become more complex and demands for power have grown, the grid has experienced several catastrophic failures that have raised concerns about our understanding of its stability and are our ability to effectively manage its complexity. Two notable recent outages include the 1996 Western States and 2003 Northeast blackout. The Northeast blackout on August 14, 2003 affected over 50 million people and has been estimated to have had an economic impact between \$4 billion and \$10 billion in the United States alone³. These and other outages have raised questions about the robustness of the grid to perturbations, and specifically whether its increased complexity has introduced system effects that make it

² These computerized monitoring and management systems are called SCADA (Supervisory Control and Data Acquisition) or control systems.

³Electric Consumers Resource Council (ELCON). The economic impacts of the August 2003 blackout. 2004.

prone to cascading failures. Recent research has suggested that this may be the case, but as will be discussed further in Section 3 of this report, we believe this question remains open.

1.2 Research Approach and Objectives

A primary objective of this project was to develop an analysis approach that could assess and ultimately help manage risks in complex infrastructures. To this end, we gave considerable attention early in the project to the formulation of a holistic modeling framework that would capture relevant system features of a complex infrastructure while simplifying analysis. To direct and ground our research, the bulk power grid was selected for focused study. Commonly applied in complex systems analysis, a node-and-link network representation was chosen as the basis for system idealization. Using this approach, a modeling framework was developed by decomposing the grid into a number of coupled network systems as shown for neighboring Utility A and Utility B in Figure 1. The physical grid, including transmission lines, load busses, and generators, is shown as the bottom layer. As will be discussed further in this report, an AC model must be used to properly simulate power flow in this layer in order to resolve system state at a sufficient level of detail for the coupled control overlay model. The consumer overlay captures the consequences associated with power loss. Unlike all previous work that we have examined that limits consideration of adverse impacts to revenue loss by the utility, we include impacts from the perspective of the consumer (e.g., the extended loss of power to a hospital may lead to loss of life)¹. The consumer overlay also drives power demand at load busses. The hierarchical mix of automated local component control (e.g., relays) and broader automated and human-in-the-loop grid management systems are portrayed as the control layer. Finally, the buying and selling of wholesale power between utilities is shown as the market layer. Since deregulation the practice of trading power between geographically remote utilities has become commonplace, creating large and unplanned-for stresses on the connecting grid infrastructure. Because of the increased inter-utility power flow, wholesale power trading practices engender potentially destabilizing conditions and have been attributed as a precipitating influence in the 1996 Western States and 2003 Northeast blackouts. Resolution of market forces is therefore critical in any comprehensive modeling framework.

The remainder of this report documents the development, implementation, and application of models for the physical, control, and market components of the bulk power grid. Section 2 describes the implementation and verification of a coupled physical-control simulator to approximate load lost from perturbations of the bulk grid (e.g., losing a transmission line). This model was systematically applied to try to identify mechanisms for cascade failure in the bulk grid, with results from this study presented in Section 3. The formulation and evaluation of an agent-based market model are discussed in Section 4. Conclusions and suggestions for future work are summarized in Section 5.

Figure 1. Conceptual model of the bulk power grid.

2. PHYSICAL GRID MODEL DEVELOPMENT

This section describes the development process for the physical grid model simulation software developed as part of this research.

2.1 Background

The goal of the modeling task was to develop a simulation tool capable of modeling a physical process (in this case, a bulk electric power system), a control process (such as a SCADA/EMS system), and a market process model. The intent was to be able to model each one of these “layers” independent of one another, as well as integrated with one another. The bulk electric power system is typically considered to be that part of the system consisting of transmission lines rated at or above 69kV, thus not including power distribution systems found in city neighborhoods. When modeling the bulk electric power system, engineers will also include major generation and load points⁴ connected to the system as part of the model. These models are used to calculate system parameters using different types of solvers, such as DC solvers, AC steady-state solvers, and even AC transient solvers, with each solver becoming more complex and requiring more computing resources and system data. At the physical process level, it was assumed that we were dealing with a lossless system (line resistance equal to zero), and that the system slack bus (necessary for steady-state simulation) was not a real bus in the system, but instead a fictitious bus with no operating limits added to the system for simulation purposes only⁵.

2.2 Theory

Common techniques for solving steady-state solutions to power systems are used in the simulation tool, such as DC power flow^{6,5}, Newton-Raphson AC power flow^{3,7,5}, and Decoupled Newton AC power flow^{3,4,5}. Also implemented and used in the simulation tool is a common method for calculating line outage distribution factors (LODF)⁸ for a power system. One area of new development in the field of steady-state solutions to power systems researched as part of this project task is approximate disturbance modeling using steady-state power flow information⁹. This work attempts to calculate sub-transient and transient data points from steady-state system data, initially focusing on generator sub-transient and transient responses to line switching events. The preliminary developments of this research have been implemented as a solver as well, appropriately named the sub-transient solver.

⁴ One example of a load point would be a city neighborhood, and in some cases an entire city.

⁵ In retrospect, this should not have been the case, as it is important to place practical limits on the system slack bus as well in order to get quality, solvable systems.

⁶ Lynn Powell, *Power System Load Flow Analysis*, New York: McGraw-Hill, 2005.

⁷ J. Duncan Glover and Mulukutla S. Sarma, *Power System Analysis and Design*, Pacific Grove: Brooks/Cole, 2002.

⁸ Bruce F. Wollenberg and Allen J. Wood, *Power Generation Operation and Control*, New York: John Wiley & Sons, 2005.

⁹ Satish Ranade, New Mexico State University, private communication, 2007.

This project task attempted to model a very simplistic SCADA/protection system available for attachment to power systems being modeled, including such devices as circuit breakers, sensors, controllers, and relays. The SCADA/protection system modeled also included an algorithm for automatic generation control (AGC). This AGC algorithm is not traditional in the sense that it monitors tie line flows between utilities to control generators, but rather it monitors system slack generation output to control generators in order to replicate the built-in functions of generators to follow changing load conditions in a power system.

2.3 Software Architecture and Implementation

The software code was written in the Java programming language in true object-oriented programming (OOP) fashion, and was designed to be modular. Our power flow solver package JPowerFlow¹⁰, for example, relies only on interfaces, making it possible for someone to develop their own representation of power system devices. As long as these devices implement the required solver interfaces, the associated systems will qualify for solving via the steady-state solvers available in our power flow solver package. In similar fashion, the SCADA package also relies on interfaces, with some basic implementations of these interfaces available. Basic implementations include circuit breakers, sensors, and controllers. More sophisticated implementations of SCADA interfaces include a load-following AGC system object and various relay protection objects (the relay interface resides in the SCADA package, but the relay implementations designed specifically for this simulator reside in the simulator package). The simulator package itself (which relies heavily on the JPowerFlow and SCADA packages, as well as the Common package, described below) is also very modular. For example, various perturb algorithms can be used to perturb parts of the system as long as they implement the Perturb interface. Another example is the simulator interface, which contains code to step through a simulation. In this sense, a simulation includes executing the solver and optionally the AGC and/or relay protection. It's possible to extend the simulator interface, overriding the step function to implement this sequence of events in any desired way. In our case, we have a DC simulator class that utilizes the DC solver in the step, and an AC simulator class that utilizes both the Newton-Raphson and the Decoupled Newton solver in the step, along with the AGC and protection code.

The Common package contains code developed by the NISAC group as part of the Loki toolkit¹¹. It provides common utilities for working with graphs (networks), and this project utilizes such classes as the generic Graph class and available graph search algorithm classes such as breadth-first search.

The basic SCADA devices (circuit breaker, sensor, and controller) are interfaces that can be implemented by power system devices themselves or can be implemented by objects that utilize reflection. For example, all load objects in our simulator implement the circuit breaker interface, which requires a trip and close method. When the trip method is called the load objects set their

¹⁰ An open source software package, copyrighted by Sandia, available at <http://www.sourceforge.net/projects/jpowerflow>

¹¹ Loki network toolkit developed originally under NISAC, a joint Sandia-Los Alamos program funded by DHS. See <http://www.sandia.gov/mission/homeland/factsheets/nisac/nisac-program-factsheet.pdf>

active variable to false, and when the close method is called they set their active variable to true. However, it is not always possible for objects to implement the sensor and controller interface, because there may be more than one component of an object that needs to be monitored and/or controlled. For example, in the case of a generator object, it may be necessary to monitor multiple variables, such as output voltage, real and reactive power output, etc. Since the sensor interface only contains the `getReading` method, there is no way to implement this method for each variable in a generator that needs to be monitored. Instead, we implement the sensor interface with an object that uses reflection to read the variable. Reflection makes it possible to call upon a method within an object via code at runtime. So, we simply tell the implemented sensor object what object we're interested in and what method were interested in (such as `getVoltage`), and it's able to get the reading for us.

The load-following AGC system object is a sophisticated object warranting more discussion. As mentioned before, this object was developed to help generators follow system load as it changes in a power system being solved in a steady-state solver. It does this by monitoring the given system's slack generators (generators connected to a system slack bus) after each steady-state solution of the given system and distributing power as needed among system generators. In a steady-state solver, the system slack bus picks up the mismatch of power between system generation and load. This power mismatch can be positive (in which case there is more load than generation) or negative (in which case there is more generation than load). Thus, after each steady-state solution, the load-following AGC system attempts to distribute the power present at the system slack bus among the other generators in the system in order to bring the system power mismatch to zero.

The various relay implementations created for our simulator are fairly sophisticated as well, and they too warrant more discussion. Four different types of relays are used in this simulator, with two relying heavily on the use of an AC solver (Newton-Raphson or Decoupled Newton), and the other two relying heavily on the use of the sub-transient solver (which in itself requires an AC solver). These relays all utilize circuit breakers, sensors, and controllers to read/write data from/to, making it possible to use these relays with any power system, not just the one implemented in this simulator. The four relays can be described as follows:

- **Voltage Protection Relay:** This relay is present on every bus in the system, and monitors voltage magnitude values for it's respective bus. Depending on the level of protection the simulation is being run at, this relay will try to keep the bus voltage magnitude within a preset range by the following means:
 - Switch on/off capacitor and reactor banks present at the bus, if any
 - Adjust output voltage of nearby generators
 - Shed optimal loads (these loads may be connected to other buses)
 - Shed local loads (loads connected to this bus)
 - Disable the bus entirely

The relay attempts to regulate voltage magnitude in this order (provided each means of regulation is available), and may try each means of protection more than once before moving on.

- **Line Flow Protection Relay:** This relay is present on every line in the system, and monitors line flow values for its respective bus. Depending on the level of protection the

simulation is being run at, this relay will try to keep the line flow below a preset limit (calculated using LODF) by the following means:

- Adjust real power output of nearby generators
- Shed optimal loads (these loads may be connected to a bus that is not an endpoint for this line)¹²
- Disable the line entirely

As with the voltage protection relay, this relay attempts to regulate line flow in this order, and may try each means of protection more than once before moving on.

- Reverse Power Relay: This relay is present on every generation bus in the system, and monitors for reverse power (power being consumed by a generation bus rather than being generated). If reverse power is sensed, this relay will begin tripping generators present on this bus offline.
- Instantaneous Overcurrent Relay: This relay is present on every generator in the system, and monitors for current values above a preset limit for the generator. If overcurrent values are sensed, the generator is tripped offline.

The voltage and line flow protection relays are configured using results from a preliminary AC solve; more specifically, the Jacobian matrix plays an important role in configuring the relays. The Jacobian matrix is made up of sensitivity values for such relationships as real power and voltage angle (tight relationship), real power and voltage magnitude (loose relationship), reactive power and voltage angle (loose relationship), and reactive power and voltage magnitude (tight relationship). These sensitivity values help to determine 1) how large of a capacitor/reactor bank to place at a bus, 2) which loads are optimal for shedding to help maintain voltage magnitude at a bus, 3) which loads are optimal for shedding to help maintain flow on a line, and 4) which generators are optimal for adjustment to help maintain flow on a line.

Preset line limits used by the line flow protection relay are calculated using Line Outage Distribution Factors (LODF). LODF analysis provides post-contingency line flow values based on line parameters (reactance, mainly) and the pre-contingency line flow value of the outage line. The distribution factors provide a quick way of determining the line flow of any line after an N-1 contingency, and are used here to calculate the rated flow of each line such that any line can handle an N-1 contingency anywhere in the system.

Reverse power and instantaneous overcurrent are events that occur as a power system moves from one running state to the next and are caused by generators dynamically reacting to frequency changes in the power system. These types of events are not evident in a steady-state simulation due to the fact that steady-state simulations do not include a notion of time or frequency. However, by utilizing current research being done in the area of approximate disturbance modeling using steady-state power flow information⁶, it is possible to calculate sub-transient power values for generators in response to line switching events. These sub-transient power values can then be used as sensor inputs to the reverse power and instantaneous overcurrent relays, which will use the power values to determine if a generator is actually

¹² We should also attempt to shed loads connected to endpoints of the line as a last resort before disabling the line, similar to what we do for buses.

consuming power (reverse power) or if a generator is creating too much current (power values are used to calculate current values).

2.4 Verification and Validation

The steady-state power flow solvers used in the simulation tool were written at Sandia by Sandia personnel. Thus, it was necessary to verify that the results produced by the solvers were correct. One way this was done was using IEEE test cases¹³ available as an IEEE Common Data Format (CDF) file. There are many test cases, each varying in the number of buses that make up the power system. Each of these test case files contains realistic and accurate (solved for) voltage and power values for the system. To verify solver results, these test cases were reset to a flat start and solved. Resetting to a flat start consists of setting all non-generation bus voltage magnitude values to 1.0 per-unit¹⁴, setting all voltage angle values to zero, and setting all generator reactive power outputs to zero. For each test case, the steady-state power flow solvers present in the JPowerFlow package generated results similar to the data present in the original test cases (before flat start), with a negligible amount of error for some of the values. Another way verification of the results was done was by comparing the results of the steady-state power flow solvers present in the JPowerFlow package to the results generated by a COTS power flow solver called PowerWorld. As before, the steady-state power flow solvers present in the JPowerFlow package generated results similar to the results generated by PowerWorld, with a negligible amount of error for some of the values.

Validation of the simulation tool developed is a much more difficult task that at this point has not been fully executed. This is mainly due to the fact that, from what we understand, this is the first time anyone has attempted to model power and control systems together in this way. However, we are confident that the control system elements are contributing to realistic results, with the results being systems with generation redistributed, capacitor banks activated, lines taken out of service, and in some cases load shed after a system perturbation such as the loss of a line. Our confidence in these results stems from conversations we've had with members of the academic community and electric utility personnel about how real systems react to disturbances today.

2.5 Usage

The physical grid model simulation tool was developed with the idea that many random, statistically accurate power systems would be modeled rather than a single real system. As discussed in the next section, this is a standard way to determine the statistical behavior of failures. The random power systems are created using degree-degree distribution data as input. Ways of obtaining this data is also discussed in the next section. Once a power system is created, loads are added to the system at each 1-degree bus. The size of each load object is specified as an input variable to the simulator in terms of megawatts (MW), along with the percentage of reactive power (MVAR) in terms of real power (MW). The total amount of system load (in MW) is kept track of and is used as input when adding generators to the system. Generators are added to the system at 2-, 3-, 4-, and 5-degree buses, and the maximum output of

¹³ Available at the University of Washington's College of Engineering website. See <http://www.ee.washington.edu/research/pstca/>

¹⁴ The per-unit system normalizes voltages, currents, impedances, and powers as percentages of pre-defined base quantities. A major advantage of per-unit is the elimination of transformer turns ratios.

each generator is also specified as an input variable to the simulator in terms of megawatts (MW). Initially, each generator's output is set at 80% of its maximum. Next, the line parameters for each transmission line in the system are configured. Reactance values are randomly generated between .001-.1, and charging values are determined using a surge impedance formula that includes reactance and line voltage rating values. Once this is done, a power system exists that can now be solved using any one of the solvers available in the JPowerFlow software package. The resulting power system is solved using the Newton-Raphson AC steady-state solver, and the solver results are used as input for configuring the system's protection. Other inputs used for configuring protection are provided as input to the simulator and include maximum and minimum per-unit voltage values at each bus (used to configure bus protection), default minimum line limit (used to configure branch protection), and the level of protection to configure. The amount of protection configured for each bus and branch in the system is based on the protection level specified, and includes such things as capacitor banks for voltage control, generator voltage adjustments for voltage control, generator power output adjustments for line flow control, and load shedding (both optimal and local) for voltage and line flow control. Once the system's protection has been configured, the protection is executed and the power system is solved again using a combination of the Newton-Raphson and Decoupled Newton AC solvers (the Decoupled Newton solver is used first, and if it fails the Newton-Raphson solver is used). This cycle is continued until no protection events occur on the power system.

The power system is now considered fully configured, and is ready to be perturbed. Inputs to the simulator for this stage of the simulation include the perturb object to be used, the number of perturbations to be performed on each power system, and the simulator object to be used. Custom perturb objects can be written to perturb the system in different ways, such as taking out the most heavily loaded transmission line in the system or taking out the bus in the system with the most generation. Custom simulator objects can also be written to step through the simulation in certain ways, such as using a DC solver or an AC solver, or using a particular load-following AGC system. A simulation step can consist of anything, and as an example one of our simulator object's step function consists of solving the power system using the Decoupled Newton AC solver, executing system protection, and distributing system generation using the load-following AGC system. As the system is perturbed and/or system protection is executed, the system may fragment into two or more systems. Each fragment created is treated as a separate power system and is subsequently run through the simulator.

As the simulator executes, changes to the power system are kept track of by each of the power system objects. For example, if a transmission line is de-energized due to a perturbation or due to protection, a flag is set within the branch object itself. Once the simulation of a system is complete, the system is analyzed to determine the results of the simulation. The results data include the total number of buses and branches in the power system, the number of buses and branches disabled during the simulation, and the total number of load that was shed (in terms of MW) due to events that occurred during the simulation.

3. ELECTRIC TRANSMISSION GRID VULNERABILITY TO CASCADE FAILURES

Applying the simulator described in the previous section, the objective of this study was to identify (or eliminate) mechanisms for cascade failure in the electric power transmission grid (hereafter “the grid”). The August 14, 2003 East Coast Blackout is often given as an example of a cascade failure^{15,16,17,18,19,20}. A cascade failure is (or should be) distinguished from a blackout caused by a large scale degradation of its components (e.g., from a 100-mile wide hurricane or a fuel shortage) in that it is supposed to start with a small isolated failure (e.g., a ground fault on a single line) that successively induces other isolated failures that “cascade” into an uncontrolled failure of the whole system; metaphors of dominoes or avalanches are often invoked^{16,17}. The cascade failure necessarily would be viewed as a system effect, caused by the way that otherwise nominal components of a system interact. In the case of a system effect, monitoring or repairing individual components wouldn’t be enough to warn of or prevent system failure; instead, the system itself would need to be rewired to achieve those goals, a much more demanding task than simply fixing individual components.

In connection with risk analysis goals of this project, we were interested in cascade failures because the risk analysis is much more complicated if one small event with moderate probability can trigger a much larger event (e.g., superblackout). In particular we are interested in the statistics of cascade failure on the grid. One of the conclusions of simple congestion models abundant in the physics literature^{19,20} is that the tail of the distribution density of power lost is a

¹⁵ M. Amin, P.F. Shewe, *Sci. Am.*, (May 2007) 60; P.F. Shewe, *The Grid* (Joseph Henry Press, 2007); S. Robinson, *SIAM News* 36 (December 2003)

<http://www.siam.org/news/news.php?id=377>

¹⁶ P. Fairley, *IEEE Spectrum*, (August 2004) 22.

¹⁷ B. A. Carreras, D. E. Newman, I. Dobson, A. B. Poole, *IEEE Transactions on Circuits and Systems I* 51 (2004) 1733. The work employs data attributed to a NERC report that we haven’t been able to obtain.

¹⁸ U.S.-Canada Power System Outage Task Force, *Final Report* (April 2004)

<http://reports.energy.gov/>.

¹⁹ B.A. Carreras, V.E. Lynch, D.E. Newman, I. Dobson, 37th Hawaii International Conference on System Sciences, Hawaii, January 2004. This is the only group we know of who employed any DC power flow calculations; all others (including their first paper, and our previous work, see Ref. 20) employed surrogates for the flow (e.g., “betweenness”). See also Prof. Dobson’s web site at <http://eceserv0.ece.wisc.edu/%7Edobson/PAPERS/complexsystemsresearch.html> for a large list of his and other related work on modeling blackouts that all argue for power-law distribution of power failure and system effects causes.

²⁰ M. L. Sachtjen, B. A. Carreras, V. E. Lynch, *Phys. Rev. E* 61 (2000) 4877; R. Albert, J. Hawoong, A.-L. Barabasi, *Nature* 406 (2000) 378; R. Albert, I. Albert, G.L. Nakarado, *Phys. Rev. E* 69 (2004) 025103(R); R. Kinney, P. Crucitti, R. Albert, V. Latora, V. arXiv:cond-mat/0410318 (2004); P. Crucitti, V. Latora, M. Marchiori, *Phys. Rev. E* 69 (2004) 045104(R); A. E. Motter, *Phys. Rev. Lett.* 93 (2004) 098701; L. da Fontoura Costa, *Phys. Rev. E* 69 (2004) 066127; J. Xu, X. F. Wang, *Physica A* 349 (2005) 685; V. Latora, M. Marchiori, *Phys. Rev. E* 71

“fat” algebraic (or “power law”²¹) rather than the expected exponential¹⁷. The consequence of a fat algebraic tail is that massive power failures are much more likely than would be predicted if the tail decayed exponentially (see Figure 2). A distribution with exponential tails follows directly from the Central Limit Theorem if one assumes random failures that are identically and independently distributed. The appearance of an algebraic tail on the other hand suggests that large failures are at the very least strongly correlated. Most risk analysis is justifiably done with the assumptions of the Central Limit Theorem but if large power failures follow a fat-tailed distribution we would need to revise risk analysis for the grid. Data on power failures has been employed to support claims of a fat tail^{16,17} but remains both sparse and uncertain. Visual inspection has been complicated by the frequent redaction of the data and their plots (Figure 2, Figure 3).

Figure 2. Interpretations of power outage distribution. The left panel (Ref. 16) and right panel (Ref. 17) correspond to different representations of the same data. The “expected outages” curve on the left panel corresponds to a distribution with an exponential tail while the “actual outages” curve corresponds to a distribution with a power-law tail. It’s less clear that a meaningful exponent of a power law could be extracted from the representation in the right panel, nevertheless Ref. 17 argues that this is evidence for a power-law distribution of power failure. It isn’t clear that “customers affected” should be proportional to “power lost”.

(2005) 015103(R); R. A. LaViolette, W. E. Beyeler, R. J. Glass, K. L. Stamber, H. Link, *Physica A* 368 (2006) 287.

²¹ That is, for large x the distribution density decays as $x^{-\alpha}$ for $1 < \alpha \leq 3$. A density with power-law tails would display a straight line with a slope of $-\alpha$ on a log-log plot (see Figure 1). For both theoretical and practical reasons, $\alpha \gg 3$ doesn’t constitute a “fat” tail, so it isn’t included in our definition of algebraic or “power law” tails

Figure 7.1. North American Power System Outages, 1984-1997

Note: The circles represent individual outages in North America between 1984 and 1997, plotted against the frequency of outages of equal or greater size over that period.

Source: Adapted from John Doyle, California Institute of Technology, "Complexity and Robustness," 1999. Data from NERC.

Figure 3. Another interpretation of power outage distribution. Another representation of the same power failure data (from Ref. 18). Here it also isn't clear that there is an exponent corresponding to a fat tail. As the original caption indicates, this is yet another redaction of the original data.

So far the data themselves haven't been enough to settle the question of algebraic vs. exponential tails²². Therefore we sought to re-examine the question with simple congestion models but with more realism than currently found in the physics literature.

3.1 The AC Electric Power (ACEP) Model

Our principal idea in this work is that the implementation of simple controls in power flow might significantly alter the distribution of cascade failures in power systems. Voltage controls are central to the control of realistic power systems. The study of cascade failure on simple models employs either DC steady-state power flow calculations^{6,19} or even simpler surrogates for the power (e.g., "betweenness"²⁰), without voltage controls. Unfortunately, voltage isn't calculated in any of these calculations but remains fixed; thus the DC model has no controls besides the slack bus and AGC. In order to implement voltage controls, the next level of realism in the description of these otherwise simple models, we were compelled to solve the non-linear AC steady-state power flow equations (instead of the much simpler, linear DC flow model) that

²² An unpublished report by A. Holmberg and S. Molin ("Using Disturbance Data to Assess Vulnerability of Electric Power Systems", Nov. 8, 2004) claims to discover a power-law distribution of power lost in Swedish (mostly Stockholm) utility data. Like the elusive NERC data, we haven't been able to recover the original data for our analysis.

provide both variable power flow and voltage. The AC power flow calculation is much more complicated than the DC calculation and the DC calculation provides a good approximation to the AC power flow; this explains why AC calculations haven't been employed in physics models of cascade failure. Unless voltage controls are required, there's no strong reason to prefer an AC calculation to DC.

Our AC Electric Power (ACEP) model (which is described in full in Section 2) obtains the steady-state AC power flow on an ensemble of constrained, controlled random networks. We constructed¹¹ the ensemble of random networks so that they were constrained to reproduce both the empirical (univariate) degree distribution²³ and the empirical (bivariate) degree-degree correlation of a prototype network²⁴, so that up to second-order statistics in the degree, each random network is a statistical representative (or surrogate) of the prototype. We chose for our prototype the transmission grid, i.e., that part of the grid consisting of lines rated at or above 69KV, of the Western Electricity Coordinating Council²⁵ (WECC, formerly the "Western Interconnect"). The resulting prototype is a sparse dissortative network with average degree of about 2.6 (Figure 4).

On each network we assigned loads to all of the one-degree nodes, of up to about 10GW. Thus up to 30% of all nodes became load nodes. The reactive load was assigned to be fraction of the real load, drawn uniformly random on the interval $[0.05, 0.15]$ so that on the reactive load was 10% of the real load. Generation was assigned to a random subset of 2-, 3-, and 4-degree nodes such that no pair of generators shared the same edge and that only about 20% of all nodes were generation nodes. The generators themselves were divided into two classes: large (rated at 155MW) and small (rated at 20MW), in ratio about 2:1. The generators are adjusted so that their initial real output was 80% of their rating. The remaining 50% of the nodes were simply interconnecting nodes with neither load nor generation.

The lines were assigned random reactance values drawn from a uniform distribution on $[0.001, 0.1]$ (per unit). Lines were protected by relays that trip the line if, in the worst case all other protective measures fail, the power flow exceeds a limit determined by a Line Outage Distribution Factors (LODF)⁸ calculation that calculates post-contingency line flows using line parameter values and pre-contingency line flows of the outage line. In assigning line limits we assumed that all lines were rated at 345KV (the WECC transmission grid employs AC lines rated from 69KV to 345KV; higher rated lines are usually DC). This assignment reflects real world practice, in contrast to the expedient found in the simple models widely published in the physics literature^{19,20} wherein one sets the limit to some fixed ratio of the initial power flow found for the network.

²³ The degree is defined to be the number of edges (lines) connected to a node (bus).

²⁴ M. Boguna and R. Pastore-Satorras, Phys. Rev. E 68 (2003) 036112.

²⁵ <http://www.wecc.biz>

Figure 4. WECC transmission grid from anonymized 2005 data²⁶. Only lines rated at 69KV or higher are shown. Node and edge locations were chosen only to provide a visually clean layout and have no geographical significance.

The assignment of voltage protection distinguishes this ACEP model from previous work in the physics literature^{19,20}. We assigned the voltage protection to each network in the ensemble as follows. We began by calculating the voltage on each of the unprotected networks and noting which buses had voltages outside of acceptable limits²⁷. Since we match generation to load, all out-of-limit voltages are initially undervoltages²⁸. We assigned capacitor banks (5 capacitors per bank) to those buses with undervoltages, determining capacitor sizes using sensitivity data from AC steady-state simulation results. The banks trip, if ever, sequentially (from smallest to largest) in order to raise the voltage to within acceptable limits. If for any bus all of its capacitors tripped without restoring that bus back to normal voltage, the entire network was rejected as unfit; a network was accepted only if all the buses were found with normal voltage and with some protection remaining. This is consistent with practice in the real world: only buses that require protection get protection, for the penalty for assigning unneeded protection is not only the cost of the hardware but also the loads that such protection itself places on the system. Many networks

²⁶ From the FAIT tool developed under NISAC, a joint Sandia-Los Alamos program funded by DHS. See http://www.sandia.gov/mission/homeland/factsheets/nisac/FAIT_factsheet.pdf

²⁷ We employed $\pm 3\%$ for the acceptable voltage range, consistent with real world practice.

²⁸ Overvoltages are rare in practice and occur here only when a fragement is formed with more generation than load. We could have but didn't protect against this rare event.

were rejected for this reason. We also rejected proposed networks either because the nonlinear AC power equations didn't converge in a reasonable number of iterations (via Newton-Raphson), or because the in the process of building the network it fragmented (due either to line failure or bus failure). This is also consistent with real world practice because operators run feasibility studies to ensure that they have planned or built reasonable networks. Nevertheless this is inconsistent with the simple models widely published in the physics literature^{19,20}, where no fitness tests are applied.

The ACEP is a steady-state model so we didn't (and couldn't) supply protection against power, voltage, or frequency transients. Frequency plays no role in this calculation even though frequency transients are known to be an important source of failure in power grids. We will resume this issue in the concluding discussion of this subtask. Nevertheless, all of those issues related to the dynamics of the system are also ignored in the simple congestion models in the physics literature that generate cascade failures and power-law distribution of power failures. The ACEP model is therefore intended to be viewed as a simple congestion model with the addition of a few simple controls. We know that this set of controls isn't formally minimal (i.e., "least restrictive") but we also know that this set is much smaller than the set of controls that exist in real world systems.

3.2 Results

The model consists of collecting an ensemble of constrained random networks each with loads, generation, line ratings (and relays), and voltage protection assigned in a manner that is both self-consistent and consistent with data and real world practice. As noted above many networks were rejected in the construction process. Those "fit" networks that did survive were subjected to one of two types of failure modes: the line with the highest flow was tripped ($P=1$), or the two lines with the two highest flows were simultaneously tripped ($P=2$)²⁹. We limited the networks in the ensemble to an average of 600 nodes³⁰. We considered four different power levels labeled³¹ $D=2, 3.5, 4.5$, and 5.5 , corresponding to a range of about 4, 6.5, 8, and 10 GW of load on each network, respectively. We present here results only for one level of protection (besides AGC and line relays), i.e., voltage protection. We didn't present the case of "no protection" because that case appears to correspond to the results already known in the physics literature^{19,20} of simple

²⁹ This is different from the usual contingency analyses executed by operators but we are specifically interested only in the propagation of possibly cascading failures that originate from a minimal set of worst case scenarios in an otherwise uncompromised system. We regard the study of systems with significant component degradation as routine and not falling in the category of "system effects".

³⁰ The size chosen is an order of magnitude smaller than the WECC grid but was the largest that would allow us to collect large samples in 24-hour runs because of the computational demands of the ACEP model. Furthermore we didn't study system size effects here because the calculations grow roughly as E^3 , where E is the number of edges in the network (doubling the system size requires a week of calculation for each ensemble). We made this compromise in order to discover primary effects, size effects being secondary even though such a study would be needed for future publication.

³¹ The designation comes from our assignment of 10 loads per load node, which are assigned uniformly randomly so that the average is $10 \times D$ MW per load node.

unprotected congestion models. Therefore we present here the results for eight ensembles for the four power levels and the two failure modes.

First we consider some key statistics in Table 1 for the real load lost in the two failure scenarios for various demands. The maximum load ever lost in the thousands of cases studied was about 10% of the initial load. The distribution of load lost is complicated, being multi-modal (failing the unimodal hypothesis with the Dip statistic³²) for all but one case. We note that the empirical maximum load lost is not far from the upper confidence limit for the 99.5th percentile, suggesting that the size of the ensemble is adequate to sample tail behavior. We note that the upper confidence limit for the 99.5th percentile is below 10% load lost in all cases.

Table 1. Statistics for load lost from model results. Statistics for power lost for **P** lines initially tripped, demand per load **D**, and **N** networks in the ensemble. In all cases the mean real load lost (**Mean**, in MW) is less than 3% of the initial real load. The standard deviation (**S.D.**) is included only for reference but not as a reliable measure of the width of the distribution. More reliable is the bootstrap estimate for the 99.5th percentile at 95% or better confidence. Note that in most cases the upper confidence limit for the 99.5th percentile is close to the empirical maximum (**Max**) for the ensembles. The bootstrap estimate of the **DIP statistic** demonstrates the failure of the hypothesis of a unimodal distribution except in the last case. The highest load cases (10GW) are highlighted; their histograms are shown in Figure 5.

P	D	N	Mean	S.D.	Max	Estimated 99.5th Percentile	Upper Confidence Limit for 99.5th Percentile	Confidence Level (%)	DIP Statistic	Prob > DIP
1	2	4312	35.5	33.0	271.1	178.3	197.0	96.6	0.053	<.001
1	3.5	1716	130.9	90.1	597.7	476.1	584.7	97.2	0.058	<.001
1	4.5	1988	131.8	73.1	624.1	432.6	488.3	97.0	0.059	<.001
1	5.5	2166	166.4	83.8	887.6	556.2	640.2	95.9	0.030	<.001
2	2	3958	68.9	44.9	430.0	258.5	273.1	95.8	0.042	<.001
2	3.5	2131	180.1	96.9	855.4	516.8	593.3	95.4	0.027	<.001
2	4.5	1490	253.1	103.6	816.5	661.5	710.4	97.9	0.026	<.001
2	5.5	1831	307.0	119.6	995.4	794.1	902.9	95.0	0.012	0.09

We expect that a distribution of cascade failures with a “fat” (e.g., power-law with an exponent no greater than 3) tail to have some events that are more than an order of magnitude above the mean but we found few that were only even as much as one order of magnitude above the mean. These results so far don’t appear to be consistent with an especially “fat” tail. Therefore we investigated more directly the tail behavior by fitting the last 90% of the ordered data of the two highest demand cases to a Pareto tail but found an exponent of at least 20, out of range of fat tail behavior. For reference we show histograms of the load lost for the highest demand cases in 5.

Figure 5.

³² J. A. Hartigan, P. M. Hartigan, Ann. Stat. 13 (1985) 70.

Figure 5. Distribution of load lost from model results. Frequency of the load lost (MW) for the highest demand ($D=5.5$, initial load about 10GW) with worst case $P=2$ in the left panel and $P=1$ in the right panel. The left panel shows the only case for which the null hypothesis of unimodality succeeds; the right panel shows a case that is at least bimodal. The farthest left peak in counts the number of networks with zero power lost in both cases.

Fitting tails is known to be a less robust method to determine their behavior than the Hill statistic³³, a widely accepted³⁴ way to determine exponents (if any) in the tail. The Hill statistic is

described by $\alpha = 1 + (n - n_0) \left(\sum_{j > n_0} \log \left(\frac{x_j}{x_{n_0}} \right) \right)^{-1}$ (see Ref. 21 for a discussion of the exponent α),

where the x_j are the ordered data, n_0 is taken here to be the index of the 75th percentile and n is the number of data points up to the 99.5th percentile. To infer an exponent α , the Hill statistic must be flat over a significant range. To calibrate our expectations, we display in Figure 6 the results from the Hill statistic for an exponential distribution and a Pareto distribution with an exponent of 3. The Hill statistic increases linearly throughout the range (except for the hook at the end) for the exponential distribution while it remains flat for the Pareto distribution. The Hill statistic for the worst case ($P=2$, $D=5.5$) is displayed in Figure 7, which doesn't yield an exponent at all but behaves more like the exponential case. We also note that the Hill statistic for our data begins at 5 and in the flattest region would correspond to an exponent of about 7. Consequently the Hill statistic either yields no exponent at all³⁴, or it yields an exponent that is much too high to qualify as a fat tail.

³³ B. M. Hill, Ann. Stat. 3 (1975) 1163.

³⁴ A. Clauset, C. R. Shalizi, M. E. J. Newman, arXiv:0706.1062 (2007).

Figure 6. Comparison of Hill statistic for an exponential (red) and a Pareto (blue, $\alpha=3$) distribution. Both distributions have the same mean and sample size (1000).

Figure 7. Hill statistic measured for the worst case ($P=2$, $D=5.5$) load lost for data in the range of the 75th to the 99.5th percentile.

3.3 Discussion and Conclusions

We studied a simple model of power congestion on transmission networks that is distinguished from other such models by the inclusion of voltage controls. In order to incorporate voltage controls, we were required to solve the much more demanding AC power flow equations instead of the usual (and for power flow alone, adequate) DC equations. When we tripped the highest

flow or the two highest flow lines, our results indicate in three different tests that the distribution of load lost doesn't have a fat tail, in contrast to published results for uncontrolled model networks^{19,20}. The range of load lost is too small for a fat-tail distribution and statistical tests fail the hypothesis of a power law distribution with an exponent no greater than 3. The presence of a fat tail is taken to be a sign of "complexity", i.e., system effect that amplify small failures into large ones. Although our maximum load lost of about 10% would be economically devastating to many utilities and customers, such magnitudes don't match superblackouts like the event of August 14, 2003.

Of course superblackouts have occurred and still occur throughout the world. The question here is whether these are preceded by large-scale component failures (including failures of control systems whether human or automatic) or inherent in the way these systems are wired together. If it turns out that the data do support the assertion of an empirical power-law distribution of power failures (which remains to us unclear), then the physical causes would need to be found in the presently intractable (for large systems) studies of dynamic effects which we ignored, such as transient deviations from 60Hz, or in some combination of steady-state and transient effects, because we found no evidence for the fat tail in the distribution of power failure from simple congestion models that included some realistic state-based static controls.

For future work we suggest a systematic study of (a) the actual power failure data (b) size effects (c) the role of control systems. The latter seems especially important for understanding this question. Although we chose a set of rules for constructing protection that roughly resembled real world practice, it was only one such set. Such a systematic study that varies the number and kind of protecting control systems might reveal the boundary of failure due to system effects.

4. PRICE DIVERGENCE IN A MODEL OF SHORT-TERM ELECTRIC POWER MARKET WITH UNSUPERVISED BILATERAL TRADES BETWEEN MYOPIC AGENTS

4.1 Introduction

As a complement to the research described in the previous sections, the objective of this study was to model the demands that markets place on the physical electric power delivery infrastructures. To do this we developed a model of market orders that employs autonomous agents that don't require human intervention. One of the key results of that work is that price divergence can occur simply as a result of the inability of agents to store electric power rather than due to panic or gouging.

The price evolution in markets for short-term electric power (Figure 8) is the opposite of the textbook examples of two-good market, for which any price fluctuations rapidly diminish as they settle down to an equilibrium price (Pindyck and Rubinfeld, 2001). In the short-term power market the price fluctuations increase dramatically towards the end of the trading session, resulting in both high and low prices that can be many standard deviations from the mean price. This phenomenon could be plausibly attributed to a variety of factors, e.g., the inelasticity of supply and demand or the structure of the market, that impute to the agents extensive market knowledge and strategic skills. Such assumptions would be consistent with both the recent experience of short-term power markets and with traditional economic theory (Pindyck and Rubinfeld, 2001).

Figure 8. Schematic of contracted price as a function of time to deliver power (taken from Stamber, 2001).

Therefore it may seem surprising that in this paper we consider instead the hypothesis that the divergence phenomenon might be attributable to random interactions between myopic agents (i.e., agents who know only about themselves) with simple preferences expressed as budget constraints and target quantities of demand or supply. This hypothesis was inspired by and parallels the work of (Gode and Sunder, 1993), who provided a minimal model of a double-auction market with what they termed “zero-intelligence” agents (who bid and ask random prices without regard to the state of the market or other agents) that produced a price evolution that rapidly converged to the same equilibrium price as that predicted by the traditional theory with its assumption of intelligent agents. Certainly their model did include the one agent endowed with knowledge of the market, i.e., the auctioneer, who would steer the randomly fluctuating prices towards the equilibrium by progressively narrowing the range in which the bids and asks would be accepted. In the markets of interest here, we consider only bilateral trades without the benefit of a market maker. In unpublished work (discussed in Appendix A), Axtell generalized the Gode-Sunder zero-intelligence agent model to a bilateral market; the divergence phenomenon was neither sought nor apparent even in that model. Nevertheless we propose that a simple modification of the original zero-intelligence assumptions would be sufficient to produce the divergence phenomenon in a bilateral market of myopic agents.

4.2 Background

The marketing of electric power resources for the purposes of maintaining system integrity while allowing for profit opportunities beyond the traditional limits of utility regulation have been in operation in the United States since the passage of the Public Utility Regulatory Policies Act of 1978 (Harris and Moncure, 2004). This was further enhanced by language in the Energy Policy Act of 1992 encouraging wholesale power competition. Orders (1996a; 1996b) later issued by the Federal Energy Regulatory Commission (FERC) codified this language into useable rules. Bilateral contracting among utilities served as one of a number of useful means for implementing these reforms.

This trend grew in the last decade and a half with the implementation of market structures (following on the heels of the FERC Orders) which allowed for long-term contracting (typically on a bilateral basis) as well as short-term ‘power pool’ markets designed for demand-gap-filling, typically centered on a Locational Marginal Pricing model, where bids to buy and sell power are settled by a central arbitrating body, and are based on the locations of supply and demand, transmission capacity constraints, and offered prices. Markets for Ancillary Services (e.g., spinning reserves, reactive power) have also followed this central arbitration model. The marketing of electric power and services has evolved from a wide variety of structures – some more successful than others – towards a Standardized Market Design (SMD), incorporating these market elements and others (Kiesling and Mannix, 2002; Zhou, 2003). We note that the SMD, as well as most currently developed energy markets, employ a combination of bilateral and central arbitration structures. This reflects the desire of market designers to allow for markets to thrive while taking into consideration that the reliable supply of electric power is paramount. Short-term settlements in a purely bilateral atmosphere might create opportunities for non-optimal solutions to the dispatch of power among market participants which, due to physical constraints of the operating system, could lead to an artificial lessening of reliability for the participants. The central arbitration structure is designed so that demand is met in the short term

at an equitable level of revenue or cost to the sellers or buyers, respectively, while ensuring that dispatch is resolved reliably.

Nonetheless, bilateral transactions serve the dominant share of marketed electricity (Zhou, 2003). Establishing bilateral transactions, both for the buyer and seller, minimizes ‘risk’. For the seller, the inherent ‘risk’ is that generation capacity which is unnecessary to meet local demand will go unused, and will not be needed in shorter-term markets due to market pressures from other sellers with similar exposures. For the buyer, the ‘risk’ is that unmet expected load, while able to be met with available capacity, will be done so only at an extraordinary premium. For both the buyer and the seller, bilateral transactions minimize the risk that the agreed-to transaction might be subsequently infeasible because of transmission constraints (e.g., congestion). The price element of this risk to buyer and seller alike is best expressed in the variance of the settlement price seen for electric power as a function of the time of settlement *ex ante*, as illustrated in Figure 8 above. Contracts placed well in advance of necessity are much more likely to be established at a value which has a minimal degree of variance from the long run marginal cost of operation of the facility (plus a small profit). Contracts placed nearer to the time of necessity face high variability in general, and engender risk to buyer and seller alike. Much of this is based on the level of demand relative to availability at the point of need. In period of low demand relative to supply, the typical seller will be faced with taking any price, even if at an usually low price for the period, for the sake of operating the sold capacity. In periods of high demand relative to supply, buyers are left with two alternatives: take whatever price is offered, or reduce demand (through planned outages, customer interruptions, and like actions). This situation can be exacerbated by the introduction of bidding structures inherent in the market structures that create the opportunity for high settlement prices (Hurlbut et al., 2004). Each of these transaction behaviors takes place with the full knowledge that most of the participants, both buyers and sellers, are profit-maximizing entities, responsible to shareholders, with all of the inherent risks (2002).

Bilateral contracting of short-term power (within the security requirements implemented in the designs of existing short-term markets) remains a useable structure, especially in transactions between and within areas of the North American power grid which have not yet implemented structures along the lines of the SMD. Here, little time is left to waste, as agreements must be negotiated and transmission access rights secured in a limited window of opportunity. Many of the markets following many of the aspects of the SMD have incorporated ‘price caps’ into the structure; however, for those which have not (or have placed high cap values), and for those areas following bilateral practices, documented prices for energy have been seen at extraordinarily high levels, up two orders of magnitude of the price under typical operating condition, typically for small quantities of power over short periods of time, necessary to maintain system integrity through times of peak demand (1998).

4.3 Description of the Model

We assumed that the market consisted of agents who traded bilaterally without the supervision of or input from a market maker. In the following paragraphs we describe the agents, the initial conditions, and the transaction rules of this model market. All of the quantities, budgets, or prices were taken to be unitless. The differences between this model and the classical Edgeworth-box trading paradigm are discussed in the conclusion of this section. The discussion

of the differences between our myopic agent and the classical zero-intelligence agent (Gode and Sunder, 1993) is a digression that we defer to Appendix A; the principal difference is that our agents possess multiple units of quantity and that our agent's preferences are expressed through budget constraints rather than with fixed costs or values. We emphasize that this minimal model doesn't provide, e.g., a real-time simulation of the transaction dynamics of an actual short-term power market (for which we lack the data in any case); instead, we employed it to test hypotheses about the information agents require in order to obtain the important features of such a market.

For each trading session we divided the agents into N_B buyers and N_S sellers. No agent could change roles, in contrast to traditional models (see discussion below). Each agent was endowed with multiple units of integer quantity (supply or demand, see Appendix B) chosen randomly from a discrete distribution; the lower bound on the distribution of the quantity was much larger than unity. We also endowed each agent with an initial budget B_0 chosen randomly from a continuous distribution. The buyer's budget was the amount she could spend to acquire her initial demand D_0 . The seller's budget was the amount that he was required to recover from the sale of his initial supply S_0 . In contrast to the buyer, his budget was allowed to become negative; in that regime, all sales would contribute to profit beyond his revenue target. Each agent always knew its own quantity and budget but knew nothing about the budget or quantity of the other agents. Buyers left the market if their demand was met or if their budget was spent (i.e., they weren't allowed to accumulate debt). Sellers left the market only if their supply sold out.

The initial budget distribution was characterized by an initial aggregate price elasticity parameter c (see Appendix B), which was selected so that the resulting initial budget-quantity curves could appear to be strongly exponential and inelastic (case EXP) or nearly linear and more elastic (case LIN) where they intersected. We considered these two extreme cases, even though case EXP seems to us to be the more realistic, because we were interested in the sensitivity of the results to different initial conditions. Figure 9 shows typical supply-demand curves for the initial conditions, where for each agent its initial value or cost would be interpreted as its B_0/D_0 or its B_0/S_0 , respectively. We stress that the curves in Figure 9 applied only to the initial conditions because we employed budget constraints rather than value or cost to determine the agent's maximum willingness to pay at any moment; therefore the intersection isn't meant to predict the mean price.

It wasn't necessary in this model for all buyers to meet their demand but it was convenient both for realism and reproducibility; alternative initial conditions that led to some buyers failing to meet their demand didn't qualitatively change any of the results presented below. It was enough to assume a 5% excess total supply to ensure that all buyers in all trading sessions met their demand, i.e., none left the market because their budgets expired first. Therefore only 5% of the sellers failed to sell all of their supply (consistent with our choices for supply and demand) and only 3% failed to meet their revenue goals (expressed in their budget) in all the trading sessions.

Figure 9. Supply-demand curves of the agent's initial conditions for a typical trading session. See Appendix B for the parameters. The linear curves (case LIN) are the most elastic where the two curves intersect; the exponential curves (case EXP) are also the least elastic at the intersection.

The budget constraints in conjunction with an agent's initial endowment of multiple units of quantity played a central role in these bilateral transactions. Each transaction began by randomly selecting a buyer and a seller from those remaining in the market. The buyer entered her bid and the seller entered his ask for exactly one unit of power, independently of each other, as follows. The buyer bid a random number drawn uniformly from the interval $(0, B/D]$, where B was her remaining budget and D was her remaining demand. The seller's ask was determined from one of two cases: For a positive budget, the seller's ask would be a random number drawn uniformly from the interval $[B/S, B]$, where B was his remaining budget and S was his remaining supply; otherwise, because we allowed negative seller budgets, he would ask a random number drawn uniformly from the interval $(0, B_0/S_0]$. If the seller's budget were positive but he had only one unit of supply left, his ask would become exactly B ; we chose budgets so that this case almost never occurred in practice. Those few sellers with a positive budget at the end of the trading session were almost always those with many more than one unit remaining unsold. Both of these choices for constructing bids or asks were conservative compared to the alternative (for the buyer) of bidding up to her entire budget for one unit or (for the seller) of always asking between zero and his whole budget. That alternative would result in some buyers failing to meet their demand and more sellers failing to meet their revenue goals but it wouldn't qualitatively change the results presented below. Another alternative does strongly impact the results, i.e., expressing the agent's preferences with fixed value or cost (see Appendix A) regardless of budget. We note that in the case each agent is endowed with exactly one unit of quantity the budget constraint is equivalent to fixed value or cost; nevertheless, whether the agents are endowed with one or many

units, that alternative wouldn't produce the funnel-shaped price divergence anticipated in Figure 8.

The transaction succeeded if and only if the buyer's bid exceeded the seller's ask. In that case, the sales price was determined by randomly distributing the surplus between the buyer and seller according to

$$price = (bid - ask) \cdot \kappa + ask$$

with κ selected uniformly randomly on the unit interval (alternatively, one could have fixed κ anywhere on the unit interval without qualitatively changing the results presented below). Furthermore, each agent's budget was reduced by the sales price and each agent's quantity would be reduced by one unit. The offers were presented as "take it or leave it", so that if the transaction failed, there would be no subsequent negotiation. No agent employed the history of bids or asks to calculate future bids or asks. There was no restriction on or charges for the number of transactions that were attempted; instead, buyers and sellers continued to be randomly paired until either there were no buyers left or no sellers left (it turned out that there were always sellers left because the buyers always met their demands with the parameters specified in Appendix B). A particular buyer-seller pair could be drawn randomly more than once because agents were sampled "with replacement" until an agent was removed from the market. Transactions that were agreed upon were assumed to be feasible and free from transmission charges.

Each attempted transaction (successful or not) counted as one unitless step in the trading session. The step plays the role of time only in the sense of imposing an ordering on the transactions. In a real short-term power market, both buyers and sellers are in a race against the clock. Here, the trading session was allowed to run as long as there was both supply and demand. Therefore the time pressure on the agents manifested itself exclusively through the shrinking supply and demand; agents weren't given clocks and could infer "time" only from their remaining budget and quantity.

This model differs from the classical Edgeworth-box trading paradigm (Pindyck and Rubinfeld, 2001) in three key ways: we abandoned the classical concept of traders by fixing an agent as either a buyer or a seller; we imposed a more restrictive specification of trading preferences; we allowed at most one unit of power to be "sold" per transaction without any restriction on the other good (money). In the classic two-good trading models without production, all traders are assumed to be consumers of each good who are endowed with some initial quantity. According to their indifference curves, they would attempt to make a trade which would increase (or at least not decrease) their utility. There would be no restriction on the number of units sold for each good, except that they must be feasible according to the agent's endowments. In making a trade, there would be no pre-determined buyer or seller; indeed a trader could switch between buying and selling as necessary to increase his utility.

In our model, we imposed a more rigid structure. Buyers were identified and given a target quantity to buy; once it was obtained they were satisfied (never to buy more) and wouldn't become sellers. Likewise, sellers were given a target quantity to sell; once the supply was sold, they wouldn't become a buyer. Thus, we adopted the terms "buyers" or "sellers" to represent our trader agents since we restricted their behavior accordingly. This rigid specification of buyer

or seller preferences isn't often adopted in classical economics. Typically preferences exhibit diminishing marginal utility and local nonsatiation, which were not adopted in this model. Our buyer wasn't required to value the last unit of power any less than the first. As a result, she may have paid a higher price for her last unit of good than she did for her first. Additionally, more is not always better for the buyer, who will exit the market once her target demand is met. Physically this is a consequence of the impossibility of storing bulk power. Instead of allowing for substitutability between the two goods, we assumed that buyers were only interested in satisfying their target demand and sellers were only interested in selling their target supply given their respective budget and cost constraints.

4.4 Results

For each of ten independent trading sessions we initialized two sets of buyers and sellers, one for each of the two cases EXP and LIN (see Appendix B for the parameters). In each set, each agent was assigned randomly selected budgets and quantities. The two cases produced similar results. Figure 10 shows two typical trajectories of the sales prices for the two cases, respectively. We note immediately the qualitative resemblance of both trajectories to the funnel shaped curve sketched in Figure 8 as the prices diverge both upward and downward away from the mean price in the latter part of the trajectories. The principal differences between the two trajectories are that many more transactions were attempted in case LIN to close the market than in case EXP and that maximum prices in case LIN were lower than in case EXP. For both trajectories, there was a period for which the prices fluctuate narrowly, corresponding to an elastic regime for both sellers and buyers, but the prices began to diverge after about 10000 steps, reaching their most extreme divergence near the end of the session.

Figure 10. Log-log plot of the sales price evolution generated in a typical trading session for cases EXP and LIN, respectively.

Figure 10 reproduces the typical trajectory for case EXP (but with a linear price scale) in order to show the typical rise of the maximum price, which was most dramatic as the market approached its close. We also include in Figure 11 the running mean sales price and its variance, in order to show how they track the growth of the maximum price.

Figure 11. Sales price evolution of a typical trading session. Results are for Case EXP, shown also in Figure 10, in which the evolution of the maximum price is indicated. The instantaneous price, the maximum price, the mean price (“running ave.”) are read on the left-hand axis. The variance (“running variance”) is read on the right-hand axis.

The large divergences shown in Figure 10 and Figure 11, although anticipated (in fact required for a realistic treatment of this market), complicate the statistical analysis because the trajectories never settled upon a steady-state price. Furthermore, the resulting price distributions were far from the normal, complicating the interpretation of statistics based on higher moments. We note that the mean sales price (36 for EXP, 35 for LIN) didn’t vary much throughout the trajectories despite the large and rapidly growing variance, which in turn qualitatively tracked the growth of the maximum sales price.

Figure 12 shows the superposition of the ten trajectories of the maximum price of case EXP, along with the change in the maximum price. We also note that large changes in the maximum price didn’t begin until after about the first one-fifth of each trajectory had evolved. The maximum change in the maximum price for each trajectory (colored in red) occurred with one exception in the latter third of the trajectories.

Figure 12. Maximum price evolution of the ten trajectories (i.e., trading sessions). The evolution of the change in the maximum price and the maximum (for each trajectory) of the change of maximum price are indicated with triangles.

The empirical probability functions (i.e., cumulative distribution F functions) of the sales prices from all ten trajectories are displayed in Figure; the two cases nearly overlap. We also displayed the empirical probability functions curves for the New England (NE-ISO) and California (CAL-ISO) sales price data (in USD/MWh); the NE-ISO data were collected from hourly reports for all of 1999-2002 and the CAL-ISO data were collected from April 1998 through January 2001. The comparison of the model with the data is problematic, as we discuss below, because the price data results from hourly auctions instead of bilateral trades. The “S”-shaped curve on a semi-log plot shows that all of the curves in Figure 13 resemble a “log-normal” distribution even though none of the curves strictly fit the log-normal. The probability functions are fundamental and don’t require a bin size to be chosen in advance. The density distribution function in Figure 14 requires a choice for the bin size (here, 1) but density may be more intuitive; in particular it is easier to see that the high price tail behavior is similar between the models and the data.

Figure 13. Comparison of model results to industry sales price data. Semi-log plot of the empirical cumulative distributions of the hourly sales price (USD/MWh) from NE-ISO and CAL-ISO compared with the distributions of the instantaneous (and unitless) sales price from the two cases of the model.

Figure 14. Log-log plot of the empirical density distributions of the sales price. The case LIN is omitted because it would be obscured (except for the lowest prices) by case EXP.

4.5 Discussion and Conclusions

From inspection of the sales price trajectories themselves (Figure 10 and Figure 11) and the maximum prices in each of those trajectories (Figure 11 and Figure 12) we note that the sales prices diverge in the latter part of the trajectory from the mean price found in the early part of the trajectories. The elasticity or inelasticity of the initial conditions (Figure 9) had little to do with the outcomes of the model or their comparison with data (Figure 13). With budget constraints for multiple units of supply or demand, the agents cautiously but continuously adjusted their ability to pay according to their budgets so that the results were less sensitive to initial conditions with, e.g., preferences based on fixed values and costs [G-S]. We note that the mean sales prices corresponding to the cases of elastic (LIN) and inelastic (EXP) initial conditions are similar to each other and that both are a little larger than the equilibrium price (about 30) that would be expected from a simple supply-demand curve analysis of the initial conditions; the increase was due (Figure 11) almost entirely to the divergence of prices especially in the last third of the market trajectory.

The comparison to the hourly sales price data is problematic because it records prices (USD/MWh) from auctions, not short-term bilateral trades; we would have compared with bilateral trade data if it had been publicly available. Nevertheless the comparison is useful if for no other reason than to investigate the sensitivity of price divergence to market structure. First, we note that the price distribution from data is narrower than the model. This is expected when using bilateral trading versus auctions (Pindyck and Rubinfeld, 2001). Auctions and bilateral trades are two very different market structures; nevertheless the behavior in these two structures is similar for high prices. In particular we note in Figure 14 that the tail of the price distribution decreases at roughly the same rate for the model as for the data. We made no attempt to fit or censor either the data or the results from the model displayed in Figure 13 and Figure 14; in particular, the CAL-ISO data includes prices from intervals when supply and demand might have been manipulated (Sweeney, 2002).

In our model the price divergence results from, on the one hand, the ability of some buyers to spend the remaining budget that resulted from making beneficial trades early in the market in order to obtain increasingly scarce supply to satisfy their demand and on the other hand, the ability of some sellers to sell their excess supply cheaply after they had achieved their revenue goals. For both, the preference to either meet their entire demand, or sell their entire supply motivates each to continue trading even as prices diverge. It may be that price divergence isn't often observed in other markets because those agents may be more willing to refrain from trading altogether, especially if the quantity being traded is storable. Our results suggest that price divergence is a phenomenon that is insensitive to the details of market structure and to the knowledge that agents might have about each other and of the market itself.

We conclude by observing that a model market of myopic agents, pioneered by the Gode-Sunder “zero-intelligence” model, is on the one hand extremely simple and on the other hand able to account for many qualitative features of the short-term electric power markets. In particular, the price divergence seen in such markets might be due to wide range of causes or exogenous shocks, e.g., a deep knowledge of the market, panic, artifacts of market policies, or clever strategies by the agents; nevertheless this simple model indicates that such assumptions are unnecessary.

We conclude by suggesting that future work incorporate this autonomous market model into the electric power grid simulations of the previous section, in which the market model provides the inputs to the simulation of the grid.

4.6 Appendix A

Here we digress to address the differences between what has become the usual development of zero-intelligence agents (ZIA) and our use of myopic agents. The ZIA model of a double-auction market was introduced by (Gode and Sunder, 1993) in order to show that the rules or constraints of the market overwhelmed the role of the individual trader in determining the market's equilibrium price; in particular they showed that the zero-intelligence traders could find the same equilibrium price predicted from the supply-demand curve, which by contrast assumed all-knowing agents. Gode and Sunder's work has since been widely discussed in the literature. In particular, (Cliff and Bruten, 1998) expanded this work to argue that the symmetric market supply-demand curves employed by Gode and Sunder contributed to the convergence to the equilibrium price.

In any case the one remaining intelligent agent, i.e., the auction house (or market maker) plays a crucial role in the Gode-Sunder ZIA model. On the other hand, there is no market maker in a market of agents making bilateral trades. Axtell has extended the Gode-Sunder model with an unpublished ZIA bilateral market, apparently developed as a pedagogical aid. As with Gode and Sunder, the agents have exactly one unit to buy or sell. Buyers and sellers don't switch roles. The supply-demand curve is linear and symmetric. Each randomly bids or asks on the interval $(0, B]$, where B is their value or cost, respectively, for their sole unit. A bid is accepted if and only if the bid exceeds the ask and the surplus is divided randomly exactly as discussed in the description of our model. The restriction of agents to one unit means that their budgets are identical to their value or costs. The result of this model is that prices fluctuate much more than they do in a double-auction market but these fluctuations nevertheless remain confined to the narrow value-cost bounds set by the agents. If allowed enough time, the price evolution in Axtell's version also converges (albeit more noisily) to an equilibrium price. This is portrayed in Figure A, which shows the price evolution with Axtell's code for 50000 agents equally divided between buyers and sellers.

Figure A. Price evolution with Axtell's bilateral ZIA market model. Each of 25000 agents is endowed with exactly one unit of supply or demand. Each agent randomly chooses its own value or cost B uniformly from the interval $(0, 30)$; this produces symmetric linear supply-demand curves that intersect at $P = 15$ and $Q = 12500$. A transaction is attempted by pairing a randomly drawn buyer j and a seller k ; j bids randomly from the interval $(0, B_j]$ and k asks from the interval $[B_k, 30)$. In this trading session 14171 units were transacted for a mean price of 14.9 (equilibrium price was 15.0) in 10^8 steps.

On the one hand it is gratifying that Axtell's bilateral ZIA model can also find the equilibrium price even without the auctioneer; this was the behavior that was sought. On the other hand, this convergent price evolution is the opposite of we expect in short-term power markets and what we did obtain with our myopic agents. By assigning multiple units of quantity to each of our agents and employing budget constraints, our myopic agents adjust their preferences against their remaining budget. It is as though we have collected the one-unit agents together into a firm and endowed them with the ability to compare notes and adjust their effective values or costs according to the status of the other agents in their firm. In this sense our agents violate the pure non-adaptive ZIA assumptions of Axtell and Gode and Sunder; nevertheless they remain myopic to the extent that they remain ignorant of both the market and other agents.

4.7 Appendix B

Here we present details of the initialization of the agents. We assigned each agent's initial quantity (supply S_0 or demand D_0) from the nearest integer of a random variate drawn uniformly from the interval $[lo, hi]$, where lo and hi are given in Table A below. Then we assigned each agent's budget first by drawing a random variate X from the continuous probability distribution $F(x)$, where x is a provisional cost or value, and

$$F(x) = \frac{\exp(-c \cdot lo) - \exp(-c \cdot x)}{\exp(-c \cdot lo) - \exp(-c \cdot hi)}, x \in [lo, hi],$$

with both x and c real. The initial budget B_0 for each agent was formed from the product of X and the initial quantity; therefore the initial cost or value is interpreted as B_0/D_0 or B_0/S_0 , respectively. For small $|c|$, F is essentially linear, giving rise to linear supply-demand curves (case LIN); otherwise, the supply-demand curve is essentially exponential (case EXP). For the choice of parameters listed in the Table, supply and demand for case EXP was inelastic where the supply-demand curves crossed, while supply and demand for case LIN was much more elastic (see Figure 9) in the same region. We created 10 sets of agents for each of the two cases, with $N_B = 1000$ buyers and $N_S = 500$ sellers for each set. The choice of N_B and N_S were inconsequential; reversing these numbers gave substantially the same results, as did choosing both to be the same.

Table A. Parameters employed in initializing the agents for the two cases EXP and LIN.

	EXP	LIN
Supply		
<i>lo</i>	10	10
<i>hi</i>	100	100
Cost		
<i>lo</i>	10	10
<i>hi</i>	200	35
<i>c</i>	-0.15	-0.01
Demand		
<i>lo</i>	10	10
<i>hi</i>	42	42
Value		
<i>lo</i>	5	30
<i>hi</i>	50	50
<i>c</i>	0.2	0.01

For each set, the total demand by all the buyers was about 26000 and the total supply from all the sellers was about 27500 units of power. We chose the parameters so that in both cases the supply-demand curves (see Figure 9) intersected at a price of about 30 for each set.

4.8 References

- 1996a. Order 888: Promoting Wholesale Competition Through Open Access Non-discriminatory Transmission Services by Public Utilities; Recovery of Stranded Costs by Public Utilities and Transmitting Utilities. United States Federal Energy Regulatory Commission.
- 1996b. Order 889: Open Access Same-Time Information System (formerly Real-Time Information Networks) and Standards of Conduct. United States Federal Energy Regulatory Commission.
1998. A Review of the Circumstances and Factors Which Resulted in Capacity Shortages and Price Volatility in Midwest Electricity Markets the Week of June 22, 1998, Indiana Utility Regulatory Commission Staff Report.
2002. Statement of Frank A. Wolak, Senate Committee on Governmental Affairs. United States Senate.
- Cliff, D., Bruten, J., 1998. Less than human: simple adaptive trading agents for CDA markets, in: Holly, S. (Ed.), *Computation in Economics, Finance and Engineering: Economic Systems*. Proceedings volume from the IFAC Symposium, 29 June-1 July 1998. Kidlington, UK : Elsevier Science, 2000, Cambridge, UK, pp. 117-122.
- Gode, D.K., Sunder, S., 1993. Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality. *Journal of Political Economy* 101, 119-137.
- Harris, M., Moncure, D., 2004. The Public Utility Regulatory Policies Act of 1978.
- Hurlbut, D., Rogas, K., Oren, S., 2004. Protecting the Market from “Hockey Stick” Pricing: How the Public Utility Commission of Texas is Dealing with Potential Price Gouging. *The Electricity Journal* 17, 26-33.
- Kiesling, L., Mannix, B., 2002. Standard Market Design in Wholesale Electricity Markets: Can FERC’s Proposed Structure Adapt to the Unknown? Reason Public Policy Institute.
- Pindyck, R.S., Rubinfeld, D.L., 2001. *Microeconomics*, Fifth ed. Prentice-Hall, Upper Saddle River.
- Stamber, K.L., 2001. Understanding the Power Crisis: Infrastructure Interdependencies Issues. Sandia National Laboratories.
- Sweeney, J.L., 2002. *The California Electricity Crisis*. Hoover Press, Stanford.
- Zhou, S., 2003. Comparison of Market Designs, in: Hurlbut, D., Greffe, R. (Eds.), *Market Oversight Division Report*. Public Utility Commission of Texas.

5. CONCLUSIONS

The purpose of this LDRD project was to develop the analysis tools and enabling insights that will allow us to design and build robust next-generation infrastructures that can withstand both terrorist and natural threats. Our research was grounded through the detailed study of a specific infrastructure, the bulk power grid. We formulated and implemented novel computer models of the grid that can partially resolve the coupled dynamics of its physical, control, and market components. In work documented separately, our bulk power simulator was integrated into a risk management framework to support improved decision making of mitigation options based upon the overall utility of decisions from the perspective of a decision maker¹.

New realism was achieved in the bulk power simulator developed for this project. Significant attention was given to increasing the fidelity of the simulator to be able to adequately model relevant features, such as relays, in both the physical power grid and control system overlay. These improvements increased the simulator's accuracy, enabling much more realistic and representative simulations of disturbance scenarios. Fidelity was advanced through the addition of new control (SCADA) elements into the AC model: four different types of relays - voltage protection, line flow protection, reverse power, and instantaneous overcurrent - were designed and implemented. Algorithms for managing power stability, such as automatic generator control and load shedding, were also developed and tested. Finally, in collaboration with NMSU, we also developed and tested a new approach to approximate disturbance modeling using steady-state power flow information. This advancement calculates sub-transient and transient data points from steady-state system data for generator sub-transient and transient responses to line switching events⁹.

A comprehensive study of an idealized regional power grid (a section of the WECC) was performed to better understand cascade failures. The analysis was performed through simulation of an ensemble of constrained random networks each with loads, generation, line ratings (and relays), and voltage protection assigned in a manner that was both self-consistent and consistent with data and real world practice. Our analysis is distinguished from previous studies by the inclusion of the control elements discussed above; the addition of these control elements effectively eliminated the occurrence of cascade failures in simulations. Preliminary results indicate a much rarer occurrence of large-scale blackouts than has previously been predicted.

Our analysis of grid failure mechanisms suggests that the implementation of simple controls might significantly alter the distribution of cascade failures in power systems. As discussed in Section 3, our results cast doubts on the appropriateness of the more simplistic models used recently by other researchers to examine failure phenomena in the grid from a complex systems perspective. Their simple models have generated power-law distribution of power failures; these results are cited as evidence that real large-scale power outages are emergent phenomena of a complex system, whereby one small failure can potentially lead to the widespread collapse of the grid.

The absence of cascade failures in our results raises questions about the underlying failure mechanisms responsible for widespread outages, and specifically whether these outages are due to a system effect or large-scale component degradation. This distinction is important from a

mitigation standpoint: in the case of a system effect, monitoring or repairing individual components would not be enough to warn of or prevent system failure; instead, the system itself would need to be re-engineered to achieve those goals, a much more demanding task than simply fixing individual components.

Related results from this project support other grid modernization and related critical infrastructure protection (CIP) challenges. The risk analysis methodology developed and published separately by this LDRD¹ has near-term application to utilities for compliance with the new North American Electric Reliability Council's (NERC) CIP standards, which require a systematic methodology to identify critical assets. Also, the market model developed in this LDRD hold promise as a means (when coupled to the power simulator) of evaluating the role of market forces in creating stresses on the physical grid, providing potential longer-term opportunities to support DOE with a unique modeling capability that can evaluate policy alternatives through the simulation of coupled market-grid dynamics.

6. DISTRIBUTION

2	MS9018	Central Technical Files	8944
2	MS0899	Technical Library	4536
1	MS0123	D. Chavez, LDRD Office	1011