SANDIA REPORT

SAND2009-1673
Unlimited Release
Printed February 2009

Final Report: Impacts Analysis for Cyber Attack on Electric Power Systems (National SCADA Test Bed FY08)

Jason E. Stamp, Randall A. Laviolette, Laurence R. Phillips, and Bryan T. Richardson

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-AC04-94-AL85000.

Sandia National Laboratories

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

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 ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2009-1673
Unlimited Release
Printed February 2009

Final Report: Impacts Analysis for Cyber Attack on Electric Power Systems (National SCADA Test Bed FY08)

Jason E. Stamp
Energy Systems Analysis
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1108
jestamp@sandia.gov

Randall A. Laviolette
Advanced Analytic Concepts
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1235
ralavio@sandia.gov

Laurence R. Phillips
Critical Infrastructure Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0671
lrphill@sandia.gov

Bryan T. Richardson
Critical Infrastructure Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0671
btricha@sandia.gov

Abstract

To analyze the risks due to cyber attack against control systems used in the United States electrical infrastructure, new algorithms are needed to determine the possible impacts. This research is studying the Reliability Impact of Cyber Attack (RICA) in a two-pronged approach. First, malevolent cyber actions are analyzed in terms of reduced grid reliability. Second, power system impacts are investigated using an abstraction of the grid's dynamic model. This second year of research extends the work done during the first year.

Acknowledgments

Thanks to the National SCADA Test Bed (NSTB) program at the United States Department of Energy's Office of Electricity Delivery and Energy Reliability which provided the funding for this research.

Executive Summary

The National SCADA Test Bed (NSTB) program is an effort funded by the U.S. Department of Energy (DOE) to address cyber security problems in U.S. energy infrastructure. During 2008, as part of the NSTB program, Sandia continued for the second year to investigate the electric power grid impacts that could be caused by cyber attack against grid control systems. The first year introduced the cyber-to-physical (C2P) bridge, which links cyber attack vectors to resulting events in the electric power grid (EPG), and leveraged the C2P bridge into two analysis approaches to determine grid impacts. Both approaches have been extended in this second year of work.

The first approach is Reliability Impacts from Cyber Attack (RICA). RICA combines simulation of cyber attack and a system failure/reliability model to estimate the degradation in grid performance that can be attributed to cyber attacks. The results are the estimated reduction in system availability attributable to cyber attack as measured by several indices. Our recent work has added several cyber attack scenarios and illustrated the potential for parametric analysis. Since RICA measures the performance of the grid in providing power, cyber security measures can be assessed by comparing the change in grid performance before and after the cyber security improvements. The successful operation of this framework, of which several examples are shown, allows the benefits of cyber security to be quantitatively determined as improvements in power availability. Future work will address modeling and quantifying improved system behavior brought about by cyber security measures. We are also exploring enhancing our attacker model to better represent observed behavior. Cyber attack is currently modeled in RICA as unexpected component outage at a given probabilistic rate. Although this accurately represents attacks by a class of adversary, researchers in cyber attack modeling have pointed out that it cannot accurately represent *all* adversaries.

The second approach to determining EPG impacts is based on the development of a finite state abstraction (FSA) of the infrastructure and its control systems that preserves the dynamic behavior of the system. This research continues development of the FSA analysis begun last year and has resulted in a model for a two-bus power system showing all possible discrete states and the transitions between them. The FSA research reduces the effectively infinite set of possible power grid states to a relatively tractable finite state set. The possible end states of the finite-state model, in particular the failure states, are identical to those of the effectively infinite-state dynamic power grid model. This is valuable because it enables the end states of an infinite-state grid to be determined, in theory for a grid of any size. A benefit of this work is that undesirable yet reachable grid states (e.g., widespread failure) can be found that might otherwise be seen only when and if they occur in the actual power grid. Reducing this demonstrated theory to practice for a grid of relevant size constitutes the direction of this approach.

Contents

Appendix

List of Figures

List of Tables

Chapter 1

Introduction

The impacts analysis (IA) program in the National SCADA Test Bed (NSTB) program at Sandia National Laboratories was founded in 2007 as a research approach to determine the results of cyber intrusion into electric power grid (EPG) control systems. The program continued in FY08, building on the FY07 work [1] presented at the NSTB workshop *Cyber Attacks on Control Systems: Evaluating the Real Risk*[1] and the 2008 NSTB Peer Review[2].

1.1 Research Goals

This project has explored two complementary approaches to characterizing grid impacts. In the first approach, the goal is the capability to estimate the degradation in grid reliability caused by cyber attacks. This approach is termed reliability impacts from cyber attack (RICA). RICA estimates how cyber attacks affect reliability by computing values for several reliability indices for a large power system model whose operation in simulated in two environments, one with cyber attack and one without. The *difference in reliability* between these two cases is the *grid impact* of cyber attack. The current goals of this work are to produce simulation results for the Western Electricity Coordinating Council (WECC) region and place these into a database so events of interest can be found by query. We are also examining the possibility of including more-accurate cyber attack models.

The second approach is development of power grid finite-state abstraction (FSA) models that preserve the dynamic behavior of the modeled systems. The resulting FSA is then analyzed to reveal all possible steady states that the system can achieve. We are particularly interested in failures attributable to cyber attack. This approach yields two benefits: First, it produces a provably complete set of failures states; second, it reveals failure modes that are difficult or impossible to discover within either the infinity of states of the actual system or the extremely large number of states of a dynamic model of the system. As this report is being written, FSA techniques have been demonstrated on systems of only a few elements; the current goal is to increase the size of the system to which FSA analysis applies.

[1] Held at the Albuquerque Hyatt and facilitated by Energetics, Inc. June 24, 2008. `http://www.sandia.gov/scada/workshop.htm`

[2] Held at the L'Enfant Plaza Hotel in Washington, DC and facilitated by Energetics, Inc. October 21-22, 2008. `http://www.controlsystemsroadmap.net/08nstb_peerreview.aspx`

1.2 Project Deliverables

The budget for the IA project in FY08 was $152k, and four deliverables were prescribed:

- Improved RICA analysis
- Improved FSA analysis
- Published results
- Final 2008 report

All have been accomplished; this document is the final report for 2008.

Chapter 2

Grid Reliability Impacts from Cyber Attack

2.1 Introduction

This introduction is an abridged description of the approach used to assess the reliability impacts from cyber attack (RICA). Full details are available in the 2007 project report [1].

Given a grid—an electric power system comprising transmission lines, breakers and other switches, generators, and loads—reliability is defined as the efficacy of the grid in delivering power to the loads. In the RICA approach, the average effectiveness of the power grid in meeting load demand is determined as follows: Models of individual power equipment elements are integrated into a model of the grid of interest and the load-satisfying behavior of this grid model is observed over several thousand simulated years using probabilistically determined outages for each individual power system component and empirical demand patterns. Any additional outages (e.g., line tripping due to overload), load flow, and unserved load are computed at each time step. The amount of unserved load is accumulated [2, 3] and, generally, averaged over time and reported per unit time. This approach is termed *Monte-Carlo (MC) reliability analysis* [4].

Several metrics based on unserved load and outage characteristics are computed. Overall, system reliability is measured using indices, including frequency of interruption (FOI) (in occurrences per year), Loss of Load Expectancy (LOLE) (in hours per year), loss of energy expectation (LOEE) (in MW·hr per year), duration of interruption (DOI) (in hours per interruption), energy not served per interruption (ENSI) (in MW·hr per occurrence), load curtailed per interruption (LCI) (in MW per occurrence), and energy index of reliability (EIR) (the ratio of energy served to yearly demand).

During MC simulation, unserved load is calculated once per simulated unit of time (every hour, in RICA; all times discussed in the remainder of the introduction refer to the simulation clock, not actual time). For each such calculation, each system element is independently determined to be in or out of service as follows: At the beginning of the simulation[1], for each piece of equipment, the time interval until its next failure is determined by scaling a sample from an exponential distribution by the relevant mean time to failure (MTTF). This interval is added to the current simulation time to give the item's "failure time", i.e., the time at which it will fail. All active items whose failure time is less than or equal to the current time (i.e., whose failure time has passed) are marked "Failed" and do not contribute to generation or transmission. Failed equipment returns to service

[1] Technically, outages are determined with respect to a system-wide *long sample interval* [5], a detail that need not be further considered for the purposes of this discussion.

after an idle period (during which it is ostensibly being repaired) determined by a similar sample scaled by the equipment's mean time to recover (MTTR). Each type of equipment has its own MTTF and MTTR.

Once it's been determined which elements are functioning, load flow is calculated based on the remaining transmission lines and their capacities, the structure of the remaining network (i.e., what's still connected to what), the capacities and setpoints of the remaining generators, and the loads to be served at that moment. Load magnitudes are based on empirical demand statistics. [6] calls the generation/transmission system, also referred to as the *bulk power system*, the "Hierarchical Level II (HL-II)"; HL-III includes distribution. Distribution is not included in the RICA model because aggregation of load at the substation level[2] provides sufficient resolution to develop an informative load picture at the regional and national levels.

The approach as described so far provides a measure of system reliability that accounts for random equipment failure and recovery on an item-by-item basis. We refer to such outages as *natural* to distinguish them from failures caused by cyber attack. To understand the impact of cyber attack, attacks and their effects must be modeled and added to the process of natural outages described above. Cyber attacks happen *in addition to* natural outages; both degrade grid performance, but they are represented independently in the model because they are expressed and mitigated in different ways. For each component the time until the next successful cyber attack is currently modeled using an exponentially distributed random variable that's independent of the natural outage variable and a selected mean time to attack[3] (MTTA). The MTTR for a cyber-attacked piece of equipment is based on the time required for cyber forensics, control system restoration, and device restoration.

The separate contributions of generation and transmission to whole-system reliability can be examined using RICA because both are explicitly represented. This means, for example, that RICA can be used to assess whether a cyber security budget would be better spent protecting generators or protecting transmission lines.

2.2 Reception of the FY07 RICA Work

The RICA algorithm and 2008 results have been informally well-received by academia, but an audience made up primarily of operations staff from the electric, petroleum, and pipeline industries criticized the work during a 2008 review[4]. The response from this audience indicates that we were only partially successful in explaining the work effectively. Some members of the review team expressed interest in the work, but others suggested that the project is of dubious merit. We feel this deserves a considered response, which is contained in the following section, *Justification for the Analysis Approach*.

[2]The substation marks the boundary between "distribution" and "transmission"; power transport between generation and the substation is considered *transmission*, everything below the substation is *distribution*.

[3]*Mean time to attack* is similar to *mean time to failure*, except that MTTA is the average interval between successful cyber attacks and MTTF is the average interval between random outages.

[4]The agenda and presentations are at http://events.energetics.com/v&c08, which was available at the time of this report's publication.

In particular, the utility of the Impact Analysis (IA) work was questioned. Overall, it does seem clear that some decisions will not be particularly informed by RICA-type results. The intent of RICA analysis is to quantify the reliability of a large-scale electric power grid using probabilistically represented cyber attacks and vulnerabilities and understand the impact of cyber attack on system performance. It is a research program whose purpose is to enable better-informed cyber security investment.

RICA results are intended to inform strategic decisions rather than provide tactical support. To isolate the effects of cyber security, the outage metrics are reported on an annual basis to "average out" the effects of the time of year and other local variations. We want to know, for example, that an observed change in system performance is due to cyber security measures, not, for example, the fact that it's currently winter. We note the potential for disparity over whether system performance over 10,000 simulated years with all annual and weather effects averaged out is relevant when considering day-to-day operations. We suggest RICA might appear more valuable to industry planners, investors, and other researchers deliberating among investments to improve cyber security.

The RICA assumptions also drew significant criticism from operations staff. In the *Justification for the Analysis Approach* section, below, we offer a defense of these assumptions and guidance for potential improvements grounded in the applicable literature. This defense provides evidence that we are not alone in advocating MC simulation as a way to measure the performance gain engendered by cyber security investment.

2.3 Justification for the Analysis Approach

Our approach is to quantify system behavior in the absence of malicious activity, then add malicious activity and quantify the difference using the measures described above in the *Introduction* section. The RICA approach is essentially a reliability framework; computation of these measures in the absence of adversaries is explicitly a reliability model. Unserved load is computed on an hour-by-hour basis using a steady-state power flow model of the system. At each time step, components fail (and failed components recover) at empirically derived rates[5]. RICA then attempts to supply existing loads through the remaining network from the remaining sources. Unsatisfiable loads are shed[6], and, if load still exceeds supply, additional load is shed according to a load-shedding plan. The amount of unsatisfied load, outage duration, etc. are recorded and used to compute the metrics.

In the presence of adversaries, computation proceeds as above, with the distinction that components fail at an increased rate based on a "successful attack" distribution function applied in an identical manner on a component-by-component basis. These "attack" failures are independent of, distinct from, and in addition to the "normal" failures. Three questions arise: Is a probabilistic model appropriate for evaluating cyber security? If it is, how should attackers be modeled? Finally, assuming a relevant attacker model, how should the attacks themselves be modeled?

[5] Associated secondary failures may occur due to cascading over-capacity tripping, etc.

[6] Load unsatisfiability can be caused not only by inadequate supply but also by inability to deliver power, which may occur even when available power exceeds total load.

Question 1: Are probabilistic models appropriate to evaluate cyber security?

The literature recommends knowing the *degree* to which a system is secure, as opposed to merely knowing that it *has certain features*. Several papers suggest a probabilistic approach to enable such a measure, which feature-based security assessments do not and cannot provide. The earlier work cited below poses several questions; the later work cited attempts to answer those questions. The *need* to answer these questions is not disputed by any of the cited sources.

Excerpts in this section are from analysts independently attempting to quantify the large-scale infrastructure impact of cyber attack and defense. Our intent in exhibiting these excerpts is to show that researchers attempting to measure cyber security advocate a probabilistic approach. Despite included comments concerning the benefits of this approach, our aim here is not to discuss *why* they have made this decision; we intend only to show that we are not alone in using a probability model to measure cyber security.

From Littlewood et al. [7]:

> Users are likely to be more interested in knowing the reliability of their system, expressed for example as a rate of occurrence of failures (ROCOF), than in knowing that it possesses certain structural properties or measures, or that it was developed under a particular regime. ... We hope the preceding discussion has made clear the desirability of a probability-based framework for operational security measurement.

From Soo Hoo [8]:

> An estimate of safeguard efficacy is essential to any cost-benefit calculation. Uncertainties severely complicate efforts to develop reliable measures of safeguard efficacy. Accepting the uncertainties and capturing them with probability distributions is one way to bound the problem of efficacy and prevent it from derailing the entire risk-management process.

From Madan et al. [9]:

> So far [prior to 2002], security attributes have been mostly assessed from the qualitative point of view. Qualitative evaluation of security attributes may no longer be acceptable. Instead, we need to quantify security. We propose a model for quantitative assessment of security attributes for intrusion tolerant systems based on stochastic models.

From Taylor et al. [10]:

> probability risk assessment (PRA) ... seeks to define and quantify the probability that an adverse event will occur. The benefits from performing a PRA for ... cyber attacks include numeric estimates for the allocation of security resources and an enhanced understanding of the security vulnerabilities and threats. Yet, despite the potential benefits, risk analysis for computer security has more detractors than supporters and is typically not done.

From Singh et al. [11]:

> In this paper, we present a probabilistic validation of an intrusion-tolerant replication system. The results are significant for the following reasons. First, they demonstrate the utility of probabilistic modeling for validating complex intrusion-tolerant architectures, and show that stochastic activity networks are an appropriate model representation for this purpose.

Question 2: Given that a probabilistic model is appropriate, what is the appropriate modeling framework to accommodate adversaries?

Many researchers, as discussed in this section, consider the ability to *quantify system availability* crucial to determining what kind of cyber hardening is needed to protect critical infrastructures. These researchers have examined *reliability analysis* as an approach to security analysis because it does exactly that: the prominent reliability measures are essentially system availability and outage characterization. The literature on this subject provides examples ([6], [12]) that specifically apply MC analysis to information security issues in a manner very similar to the current RICA approach. However, the analogy between reliability and security is not entirely apt, as is pointed out in some of the excerpts that follow. Note, on the other hand, the repeated statement that the probabilistic framework of reliability analysis is particularly suited for critical infrastructure cyber security.

From Littlewood et al. [7]:

> We discuss similarities between reliability and security with the intention of working towards measures of operational security similar to those that we have for reliability of systems. Very informally, these measures could involve expressions such as the rate of occurrence of security breaches (cf. rate of occurrence of failures in reliability), or the probability that a specified mission can be accomplished without a security breach (cf. reliability function).

From Jonsson and Olovsson [13]:

> Statistical tests on the data indicate that the times between consecutive breaches during the standard attack phase are exponentially distributed. This means that traditional methods for reliability modeling, e.g., Markov models, could be used.

From Taylor et al. [10]:

> Problems with cyber security assessment of risk include difficulty analyzing the risks and mitigation strategies for large complex networks and inaccuracies associated with the expected loss from security events. [Note this is precisely what our impact work is addressing] ... Our technique is currently being developed for power industry cyber security assessment and hardening [and] features self-assessment, risk estimates based on actual data, and quantifiable inputs for decision analysis. This assessment method is particularly well suited to hardening critical infrastructure systems against cyber attack and terrorism.

From Singh et al. [11]:

> Probabilistic validation through stochastic modeling is an attractive mechanism for evaluating intrusion tolerance.

From McDermott et al. [14]:

> There are at least three ways to shift toward the new paradigm: 1) from fault-tolerance approaches toward designed faults, 2) from trusted-component approaches toward stochastic faults, and 3) increasing the expressiveness of models such as stochastic process algebra to encompass practical systems. The first approach should be adopted when coming from the field of fault tolerance ... the second approach should be the first step when coming from the security community.

From Schneidewind [15]:

> This paper includes the conditional probability of security failures given the occurrence of reliability failures. In our model, we develop what is in effect a probabilistic specification of the incidence of cyber attacks. A security intrusion and the response of an intrusion-tolerant system to an attack can be modeled as a random process. This integrated approach is particularly applicable to control systems that govern the operation of critical infrastructure systems in chemical, electrical, rail, and aviation systems.

Question 3: Given a modeling framework that accommodates adversaries, how are the attacks modeled?

Modeling security using a modified reliability model is near the current state of the art. However, this model may not be entirely appropriate for representing the attacker; despite literature that suggests an exponential attack distribution [which is the distribution used in RICA], other distributions may be more accurate. The literature also suggests that modeling the attacker with a stationary distribution may not be appropriate, and some ways of dealing with this are discussed.

As we write this, our attacker model is implicit and represents an adversary who has completed the learning and experimental phases and can carry out the functional attack at a stationary rate. This is the most effective adversary from a perspective of overall impact, so our analysis will produce conservative results. The literature poses some questions about the inclusion of an attacker in the reliability model. One approach is to use attacker submodels that can be integrated into the primary model, and several researchers ([9] [11] [13]) suggest that this submodel take the form of a Markov process. John McDermott of the Naval Research Laboratory has implemented an attacker submodel using a mechanism that reduces to a Markov process [16]. The RICA approach supports the integration of subsidiary attacker models, although greater accuracy will cost more and take longer. We are presently considering extending our attacker models to provide a more accurate understanding of how cyber security affects overall system performance. Our immediate intent is to determine whether there is a difference in outcome caused by the shift to an attacker submodel. One question of interest is how the LOEE will change if individual attacks occur less frequently but cause greater impact.

From Madan et al. [9]:

> An attacker always tries to eventually send such a system into a security-failed state. Obviously, this requires the attacker to spend time or effort. In general, this time or effort is best modeled as a random variable. Depending on the nature of an attack, this random variable may follow one of several distribution functions. In this paper, we borrow some of the common distribution functions used in the field of reliability theory. Deterministic, exponential, hyper-exponential, hypo-exponential, Weibull, gamma, and log-logistic are some of the distribution functions that make sense in the context of security an

From McDermott et al. [14]:

> [I]ntrusion-tolerance and security researchers look at faults in terms of statistically dependent events caused by the hard intruder, [while] the fault tolerance literature assumes that faults ... can be described as random variables with probability distributions. However, when considering the survivability of a system, we cannot assume that the system is susceptible to only one type of fault or the other ... we must consider the failure behaviors of both classes of faults. ... we need to consider development of models based on a combination of stochastic behavior and ... specific detailed system behavior. This kind of model can encompass both types of faults and methods of dealing with them.

From McDermott [16]:

> We present a series of models that provide an example of attack-potential-based quantitative modeling of survivability for high-consequence systems. Our examples also demonstrate the significance of getting the intruder model right. Quantitative modeling of survivability for validation or measurement of high-consequence systems should be based on detailed intruder models. Detailed aspects of the intruder's attack potential can have significant impact on the expected survivability of an approach.

2.4 Additional Discussion on the Analysis Approach

Five other issues relating to the approach itself remain to be addressed: our non-standard use of the term *reliability*; the ostensible negation of cyber impact in real life by operating in N–1 mode; lack of data characterizing the cyber attacker; the greater operational applicability of scenario analysis results vs. the hard-to-apply RICA results; and, finally, the accuracy of the attacker models. The first four caused considerable discussion among the reviewers, while the fifth represents what we think is a key discussion point for the entire reliability impacts analysis effort. Each is addressed in this section.

Several reviewers noted that our use of the term *reliability* to refer to the RICA metrics does not strictly coincide with North American Electric Reliability Corporation (NERC) usage. Our use of *reliability*, while linguistically appropriate, does differ from NERC usage. We use the term to

refer specifically to the adequacy of the electric power grid to deliver energy under steady-state conditions. NERC usage encompasses not only this meaning but also some sense of resiliency, relating to the stability and dynamic performance of the power grid, which we do not consider here.

Reviewers also correctly pointed out that, since the grid is engineered and operated in an N–1 condition, loss of load due to a single cyber failure is, at worst, rather rare[7]. The essence of our response is: we agree insofar as the phrase "a single cyber failure" refers to the failure of a single grid element due to cyber attack.

The term *single cyber failure* is not well defined. Concerns over the N–1 criterion imply that it is being interpreted to mean "the failure of a single grid element due to cyber attack". The authors agree that in this sense, a single cyber failure is unlikely to be of particular interest. However, other interpretations of *single cyber failure* may be more relevant. A cyber vulnerability found at one control system site is in general likely to be found at multiple sites; i.e., all those sites that utilize the vulnerability-containing program or device. The extreme ease and speed with which malware can be distributed [17] suggests that a more realistic approach would consider a "single cyber failure" as affecting a *percentage* of elements (not just a single element) at sites operating the vulnerable mechanism. A greater or lesser percentage of sites would be affected based on site security practices and on the probability that the site is running the vulnerability-containing program or device version; exact values could be discovered for real-world vulnerabilities by examining the operational infrastructure. A national map could illustrate the installation footprints of operating system patch levels, commercial control system software, processor/firmware/software-containing hardware, and common ancillary software. This map would provide significant insight into the actual meaning of *single cyber failure* and enable more-accurate RICA results. If the information were adequately detailed and up to date, the exploitable target population could be found quickly for new vulnerabilities, thus enabling very specific risk assessment.

Furthermore, RICA does not enforce N–1, which logically entails RICA results that are conservative compared to what they would be if RICA *did* enforce N–1. In RICA, load is not shed unless supply is insufficient, whereas in the real world there are conditions (which RICA currently ignores) in which load must be shed preemptively in order to maintain N–1. For example, if total supply were equal to total load and no additional supply is available, a single generation outage would require that load be shed, simply because in that case there would not be enough power to satisfy all loads. By definition, N–1 is not met in this situation; to achieve N–1 when load becomes equal to supply, load must be purposely shed at that moment to prevent its being unexpectedly shed when an outage occurs. If N–1 were enforced during RICA simulation, therefore, load would occasionally be shed in cases where all load is satisfied; since outages happen randomly, this would sometimes happen in cases where there is no subsequent outage. Thus, RICA metric values are conservative with respect to N–1: load that *would* be curtailed under N–1 enforcement is *not* curtailed in RICA, so results from RICA slightly *underestimate* impacts, at least with respect to N–1. RICA could be modified to enforce the N–1 criterion. We welcome additional discussion on this topic.

[7] The *N–1 condition*, often referred to as simply "N–1", means that the grid is in a state such that the failure of any single element will not by itself cause load shedding or further failure.

The third issue has to do with the lack of relevant data needed to characterize the cyber attacker. The research team is painfully aware of the hypothetical nature of the values we have been using for mean time to attack (MTTA) and MTTR (and, for that matter, for the scenarios themselves—the hypothetical protection and SCADA attacks are also open to scrutiny). In the case of MTTR after cyber attack, it is our position that a utility employing this analysis process should be able to populate these numbers effectively based on their own cyber processes and procedures. As for the MTTA, representing the interval between successful attacks requires either an appropriate data set (of which only weakly relevant examples appear to be forthcoming) or consensus among a group of experts. The latter option, while not preferred, at least represents the process by which decisions on cyber security impacts are made now. The RICA approach accommodates this, requiring only that these experts reach consensus on a set of stimuli parameters, after which the determination of impacts is quantitative.

The fourth and final issue from the reviewers concerns the usefulness of the RICA approach as it relates to scenario analysis for cyber security. The position of interest is apparently that scenario analysis is more relevant to control systems security analysis than a RICA-style approach; specifically, "nothing actionable or ... insightful"[8] was to be found in the latter, presumably in contrast to the former. We think there's a balance to be struck, although there is considerable room to argue about what the complement to scenario analysis should be. Given the content of our RICA work and published statements by several other researchers quoted herein, we think probabilistic analysis is part of that complement. The article by Soo Hoo [8] discusses reliance on scenario analysis; we have included an extended quote on the subject because it encapsulates our thoughts on the subject perfectly:

> Scenario-analysis approaches are probably more common than any others, especially in small-to-medium sized enterprises. As its name implies, scenario analysis involves the construction of different scenarios by which computer security is compromised. Scenario analysis is often employed to dramatically illustrate how vulnerable an organization is to information attacks. For example, some consultants will, with their client's permission, hack into the client's information systems, obtain sensitive data, and provide the client with a report detailing the data stolen, how quickly it was obtained, and other particulars of the exploit. This 'red-teaming' exercise helps motivate the client to pursue better security and to provide further work for security consultants.
>
> Scenario-analysis techniques are also used to encourage broader brainstorming about computer-related risks. Some companies have small information technology risk management teams whose experts fan out to company divisions, provide facilitation services and technical expertise, and help the divisions understand their risk exposure. In this setting, scenarios are used to demonstrate the variety and severity of risks faced. The scenarios deemed most likely and of greatest severity are then used as the basis for developing a risk-mitigation strategy.
>
> The primary drawback of an exclusively scenario-analysis approach is its limited scope. Looking back at the event tree example from the common framework in Figure 2, scenarios essentially represent different paths through the tree. The danger of assessing only a few scenarios is the possibility that important paths may be missed,

[8] remark made by a reviewer during the 2008 NSTB Peer Review.

> leaving serious risks unaddressed. By narrowing the focus in this way, the analysis is made tractable, but incomplete. In addition, this approach also does nothing to encourage a better, more comprehensive data collection activity. Like the valuation-driven approaches, scenario analysis simplifies the assessment process, but in doing so runs the risk of fostering complacency as organizations, satisfied that they have addressed the specified scenario risks, are led into a potentially false sense of security.

It should be pointed out that Soo Hoo's analysis holds even when the scenarios are derived from knowledge of genuine adversaries and address serious, specific problems; in the above quote, his point is that the scenario approach can address real problems but, in the end, it can't be shown to have provided a secure system.

We think the key issue associated with analyzing impacts from cyber attack using the RICA approach is that the current RICA attack model (i.e., component cyber-outage duration simulated by random samples from exponential distributions based on MTTA and MTTR) assumes certain characteristics about the adversary that may not adequately characterize the entire attacker spectrum. RICA analysis uses exponentially distributed random variables in a memoryless stochastic process to depict the adversary as a group that can carry out repeated, successful attacks at will but is not sufficiently organized to develop high-impact scenarios. Based on continuous popular-press reports about attackers of this sort, we conclude the future control system cyber environment will almost certainly include such attackers. This implies the RICA approach is relevant as it currently stands, although we agree the attacker model would be improved if it were extended with relevant attacker types not currently modeled.

We continue to investigate modeling the so-called *hard attacker* discussed in references [14] and [16], with the intent of extending the RICA analysis to include this sort of attacker. From [14]: "We consider two kinds of intruders ... represent[ing] extremes that make our point [that a better approach to security modeling is needed] clear. One kind, hard intruders, have relatively high-value objectives, low risk aversion, high skills, and high resource levels. The other has no objective at all, low skills, low risk aversion, and the capability to attack any component at any point in its life cycle."

As this is being written, the RICA algorithm captures the second type of intruder but not the first. We agree with McDermott that, in general, the higher the consequence of failure, the more one should worry about hard intruders (the first type) as opposed to unfocused, random ones (the second type). McDermott addresses the hard intruder in terms of the intruder's *attack potential*: "Intruder work factor (e.g. mean time to accomplish an attack) is part of a good metric for survivability or security. However, intruder work factor is determined by the attack potential of an intruder. A work factor metric should be coupled with a description of intruder attack potential that determines it. Intruders are best characterized by directly defining their attack potential. Work factor or mean time to breach (MTTB)[9] can then be determined by modeling or experiment against the system of interest."

But (again from [16]) "MTTB quantification does not tell us if our system is safe from denial of service, information leaks, or something else, because it doesn't relate the survival statistic to

[9]McDermott's *mean time to breach* is similar to our *mean time to (successful) attack*.

the security architecture and protocols under attack and gives no indication of what security policy or claim is being violated by an intrusion." This is unfortunate, since MTTB is straightforward, and already accommodated by the current RICA paradigm, but does clearly imply the need for an improved attacker model.

Fortunately, the literature supplies direction for this. We find in [11], as cited by McDermott in [16], that "probabilistic models for intrusion-tolerant systems should, either explicitly or logically, include sub-models of the attacker, the intrusion-tolerance mechanism being used, the application, and the resource/privilege state of the system." This is the direction McDermott takes in [16].

Section 4 of [16] is *Modeling Survivability Using Different Intruders*. This is conceptual territory that the RICA work should occupy. In this paradigm, survivability is approximated by "availability in the presence of sponsored faults." The RICA work measures availability at high resolution in terms of lost load and is, therefore, consistent with this approach. [16] also recommends, regardless of attacker model, that the survivability calculation support both stochastic *and* sponsored faults—i.e., purposeful attack—which RICA permits.

In [16], survivability is computed for several attack/security combinations using performance evaluation process algebra (PEPA). Every PEPA model, including any we might specify to support our work, has a corresponding Markov process that can be solved to obtain the steady-state distribution of the Markov model states. The parameter values of this distribution would allow RICA to produce the metrics described above for attack/security combinations of interest.

The analysis in [16] leads to the conclusion that the hard intruder is the one that matters for high-consequence systems, and furthermore supplies a representational framework that includes the hard intruder. As this is being written, we are determining the effort required to extend the current RICA framework to include hard attackers.

2.5 Algorithm Improvement

We leave the discussion about the usefulness of the RICA approach to focus on the actual research. Last year, our results were based on cyber attack leading to generator unavailability. As stipulated, this year's work is intended to demonstrate a second-generation version of the RICA algorithm that generated last year's results. To that end, we successfully simulated the impacts to grid reliability from attacks against wind protection and system supervisory control and data acquisition (SCADA).

2.5.1 Attacks Against Protection

To study protection attacks, we focus on attacks against two common types of protective relaying:

- Generator protection
- Line protection

Although generation protection attacks are generally unlikely[10], they are included as an example since loss of generation is frequently modeled in reliability studies (although not, at least in the prior literature, from cyber attack). Based on a certain linked set of parameters from the cyber-to-physical (C2P) bridge, a successful attack will be modeled as causing the generator to go offline as its breakers trip. The interval between successful attacks is modeled using an exponentially-distributed random variable and some selected MTTA. The MTTR for the generator depends on the time required for cyber forensics and control system restoration, and the restart time for the generator, which depends on its classification (hydro, coal, nuclear, etc.).

Line protection attacks assume advanced relays that allow cyber control of the protective breakers at both ends of a transmission line. A successful cyber attack opens the breakers and removes the line from service. The line is successfully attacked at random intervals denoted by an MTTA and is out of service for an MTTR, as above. This is again similar to the process for natural outages, but with the recovery interval dependent on cyber remediation and line restoration.

2.5.2 Attacks Against SCADA

We also modeled an adversary's successful penetration and use of a grid's SCADA system to send trip signals to system breakers. Opening breakers can isolate generators, open lines, and disconnect shunts and loads. The interval between successful attacks is determined by sampling from an exponentially-distributed random variable scaled by a specified MTTA, in a manner similar that used in the protection modeling. The downtime (in this case, forced-breaker-open duration) or MTTR is the cyber forensics interval in addition to the component restoration time. Another parameter for analysis is the average percentage of tripped breakers (APTB) stemming from a successful attack, which is modeled using a Bernoulli random variable for each breaker with a selected mean value p_{SCADA}. The APTB parameter enables quantification of the concept that not all breakers will be sent a trip signal by the attack, not all breakers sent a trip signal will receive it, and not all breakers receiving a trip signal will trip.

2.5.3 Overall Component Models Including Cyber

The overall diagram for constituent grid components is shown in Figure 2.1. A generator, line, shunt, or load always starts in service, and eventually returns to service. In the meantime, random failures not caused by attack may cause it to transition to the repair state and back (This is the conventional behavior for reliability studies; note no de-rated states are currently used in RICA). A successful cyber attack against a component's protection scheme causes it to become unavailable, and an interval of protection forensics begins. This presumably results in the protection scheme being restored or a workaround developed, after which the grid element is restarted and reconnected to the grid (which may involve an additional delay, such as for a warm restart of a generator). If a SCADA attack causes de-energization of the component, then it enters the SCADA forensics interval before being reconnected.

[10]Most generators reside in generating stations, where protection significantly limits the possible attack paths.

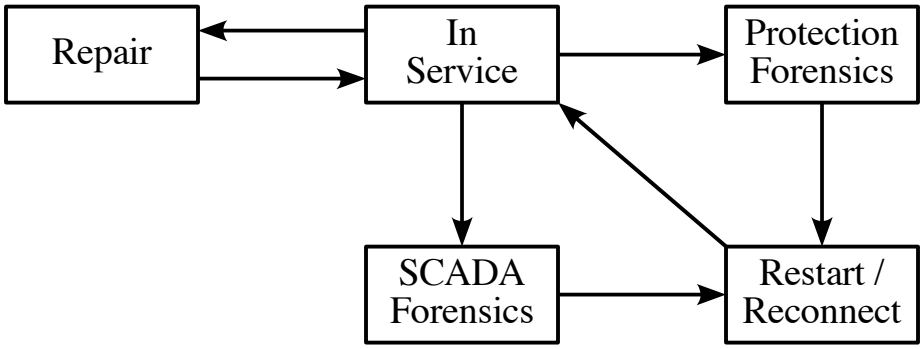

Figure 2.1: Possible states for grid elements using the RICA approach.

2.6 RICA Simulations and Results

The test system used for the initial RICA investigation into reliability impacts is the Roy Billinton test system (RBTS) [18]. This reference provides parameters for the lines, generators, and loads, including their normal failure and repair rates. The test system was modified by removing the radial line connecting bus 5 and bus 6 and adding the bus-6 load to the bus-5 load. This prevents the effects of losing this single line from dominating the simulation results [19]. This reduces the six-bus, nine-line RBTS to five buses and eight lines, although it retains the eleven generators and four loads of the original. Total system generation is 240MW and peak load is 185MW. A time-varying load model was included in the analysis, as opposed to assuming unrealistically constant peak load demand. Additional parameters for cyber forensics and recovery times are listed in Tables 2.1 through 2.3.

2.6.1 Device Restoration Delays from Cyber

For this analysis, the forensics and restart/reconnection delays were fixed (i.e., not randomly distributed) to model the likely procedural rigidity associated with control system forensics and generator restoration. Assuming the only effect of a cyber attack is equipment deactivation, the duration of restoration after a cyber attack will be less variable than after an arbitrary failure that may or may not require physical repair. The fixed forensics intervals we used are arbitrary, as there are very few data about these procedures (we found no non-anecdotal information). These values approximate delays intended to be representative of a well-established procedure, implying that an *ad hoc* or unpracticed response may cause reliability effects greater than this paper's results indicate. The generator restart times are intended to be representative of hot restart times, given the expected forensic downtime. Lines, shunts, and loads are expected to be reconnected immediately following the cyber forensics interval, although these values may be adjusted based on new information. Future research should investigate the parametric sensitivity of the RICA results to the restoration model and its parameters.

Table 2.1: Device forensics intervals for cyber attack.

Attack Vector	Interval
Generation protection	8 hours
Transmission protection	4 hours
SCADA	4 hours

Table 2.2: Restart delays for RBTS generators after cyber attack.

Unit Type	Size	Delay
Hydro	Any	+0 hours to restart
Thermal	Any	+2 hours to restart

Table 2.3: Restoration delays for grid elements after cyber attack.

Component	Delay
Lines	+0 hours to reconnect
Shunts	+0 hours to reconnect
Loads	+0 hours to reconnect

2.6.2 Test Results

The MC simulation was first used to calculate the base case measures of reliability for FOI, LOLE, and LOEE; subsequently, values for the remaining reliability indices were derived. In this formulation, the simulation uses a DC load flow approximation. Load curtailments may result from insufficient generation or transmission congestion, although they are minimized using an optimal power flow (OPF) routine to simulate remedial dispatch by system operators.

The results of RICA analysis with varying cyber scenarios are tabulated in Table 2.4. Here, λ_{GEN} is the frequency of successful attack against cyber generator protection (in occurrences per year), λ_{LINE} is that of cyber line protection, and λ_{SCADA} is for SCADA attack, while p_{SCADA} is the APTB for successful SCADA penetration.

Cyber attack against protection and SCADA (with the parameters previous mentioned) caused significantly degraded reliability for this system. In this case, the test system often incurred a higher incidence of outages. In Sim 1, the FOI indicates that interruptions are occurring more than twice as often. Given the increased likelihood of interruption, the LOLE and LOEE naturally climb. However, the drop in DOI and ENSI for Sim 1 also indicate that the additional outages caused by the protection cyber attacks are shorter than base case outages. This may be an effect particular to very reliable systems: in a less reliable grid, normal outages would occur more often, so cyber attacks and normal outages would be more likely to overlap. This could significantly increase the *amount* of curtailed load, not just *how often* load is curtailed. This may be explored in future research in this area.

Decreasing the MTTA for protection attacks (Sims 2 and 3) shows a trend of more frequent but slightly less severe interruptions. The results indicate that reliability is more sensitive to line protection MTTA than to generator protection MTTA. Given the large generation capacity margin of the RBTS test system, this is not surprising; we would expect a system with a less excess generation capacity to be more sensitive to generator interruptions. Overall, Sim 1, 2, and 3 indicate the degree to which attacks against generators and lines can affect grid reliability indices.

Sim 4 shows that even well-contained SCADA attacks (i.e., attacks that affect only 20% of the breakers, represented in the model as APTB) can significantly affect reliability indices. Even though SCADA attack frequency was lower, all the reliability indices worsened. Interruptions occurred four times as often, and expected load loss was fourteen times greater. These results are alarming in light of the fact that only one in five breakers were tripped in the successful attacks. A significant SCADA penetration could affect more than 20% of breakers, which makes the current results alarming. In Sim 5, with SCADA attacks occurring more frequently, a trend toward more severe interruption is evident, with the doubling in LOEE of particular note. Finally, combining protection and SCADA attacks causes additional reduction in reliability, as shown in Sim 6. The effect is dominated by the SCADA attack. However, the overall conclusion is that even small reliability reductions on grid components, such as those we modeled as having been caused by cyber attack, can significantly increase expected load curtailments, even in a highly reliable system.

Table 2.4: Change in RBTS reliability indices, with cyber attack.

Parameter	Base Case	Sim 1	Sim 2	Sim 3	Sim 4	Sim 5	Sim 6	Units
λ_{GEN}	0	1.0	1.0	2.5	0	0	1.0	occ/yr
λ_{LINE}	0	1.0	2.5	2.5	0	0	1.0	occ/yr
λ_{SCADA}	0	0	0	0	1.0	2.5	1.0	occ/yr
p_{SCADA}	0	0	0	0	0.2	0.2	0.2	(none)
FOI	0.048	0.112	0.161	0.169	0.742	1.819	0.809	int/yr
LOLE	0.21	0.36	0.466	0.512	2.96	7.14	3.14	hr/yr
LOEE	3.15	4.87	7.70	7.02	111.5	249.2	105.9	MW·hr/yr
DOI	4.45	3.21	2.89	3.03	3.99	3.93	3.88	hr/int
ENSI	66.2	43.5	47.6	41.5	150.3	137.0	130.9	MW·hr/int
LCI	14.9	13.6	16.5	13.7	37.7	34.9	33.7	MW/int
EIR	0.999997	0.999995	0.999992	0.999993	0.999888	0.999750	0.999894	(none)

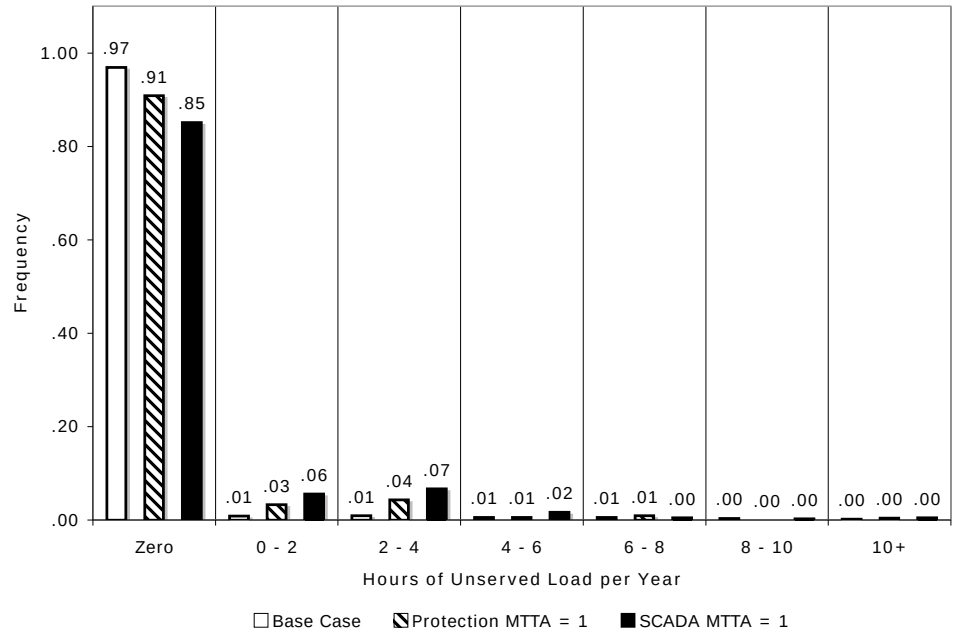

Figure 2.2: Comparing LLD (hours/year), with the base case and protection/SCADA attack.

In another interesting data comparison, the histogram for load loss duration (LLD)—the total length in hours of outages in a particular year—with cyber attack against protection (Sim 1) and SCADA (Sim 4) is compared with the results from the base case simulation in Figure 2.2. It is evident that cyber attacks as modeled cause load curtailment much more often, as the relative occurrences of years with zero hours of load lost has dropped precipitously while the rest have increased. This indicates that both the frequency and total hours of unserved load have gone up significantly when cyber attack is modeled.

2.7 Section Conclusions

Modeling the reduction in reliability that results from cyber attack allows quantitative analysis for risk reduction. Using this approach, for example, amelioration of a set of vulnerabilities could be prioritized based on the LOEE they induce. Used another way, a model of the power grid that included the cyber component could be analyzed for sensitivity to particular attack approaches, so that complementary threat analyses could be seeded with scenarios of interest. Finally, if risk reduction efforts can be represented in terms of corresponding MTTA or MTTR reduction, comparative RICA analysis can quantitatively indicate the value of the proposed mitigations. We are considering the relevance of advanced attacker models in providing informative analysis results.

Chapter 3

Grid Dynamic Impacts from Cyber Attack

This chapter details the progress made in FY08 on the finite state abstraction (FSA) modeling effort. The goal of this work is to develop a model of the power grid and its associated control systems that lends itself to analysis of cyber security attack. In particular, we are looking for opportunities for large-scale power-grid impacts that are manifest through adversarial tampering with control systems.

Last year's work (detailed in reference [1]) focused on the development of the model and its expression in a form suitable for the application of the FSA process. This year, we have successfully analyzed a two-bus model and converted its continuous-time-domain representation into a finite state system. In FY09, we will further analyze the resulting FSA of the two-bus model for cyber security vulnerabilities.

3.1 Review of the Two-Bus Test System

The model is shown in Figure 3.1.

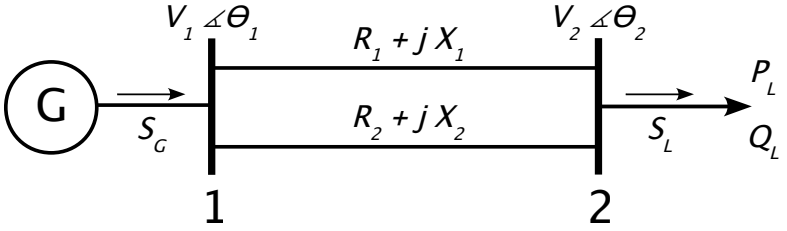

Figure 3.1: Two-bus test system.

This small-scale model encompasses many dynamic power system characteristics, including

- Generator rotor dynamics
- Machine governor control
- Voltage-based load variation
- Dynamic load recovery

Overall, the system variables include

$$x = \begin{bmatrix} \delta \\ \Delta\omega \\ P_M \\ E'_q \\ P_L \\ Q_L \end{bmatrix}, \; z = \begin{bmatrix} V_1 \\ V_2 \\ \theta_1 \\ \theta_2 \end{bmatrix}, \; u = \begin{bmatrix} E_f \end{bmatrix}, \; u_d = \begin{bmatrix} P_0(t) \\ Q_0(t) \end{bmatrix}, \text{ and } \; q = \begin{bmatrix} q_{LS} \\ q_{LT1} \\ q_{LT2} \end{bmatrix}.$$

In this system, u_d are disturbance functions representing the nominal active and reactive power demands for the load. Other variable definitions are contained in Appendix B, Table B.2. These variables fit the conventional system of dynamic equations:

$$\begin{aligned} \dot{x} &= f(x,z,u,u_d,q), \\ 0 &= g(x,q,z), \text{ and} \\ y &= Cx. \end{aligned} \tag{3.1}$$

Overall, the equations for the system dynamics are

$$\dot{x} = \begin{bmatrix} 2\pi 60\,\Delta\omega \\ \frac{1}{2H}\{P_M - P_G - D\Delta\omega\} \\ -G\Delta\omega \\ \frac{1}{T'_{d0}X'_d}\left\{-X_d E'_q + X'_d E_f + \left(X_d - X'_d\right)V_1\cos(\delta-\theta_1)\right\} \\ \frac{1}{T_L}\left\{-P_L + (1-K_{LS}\,q_{LS})\,P_0(t)\left(\frac{V_2}{V_{ref}}\right)^{\alpha}\right\} \\ \frac{1}{T_L}\left\{-Q_L + (1-K_{LS}\,q_{LS})\,Q_0(t)\left(\frac{V_2}{V_{ref}}\right)^{\beta}\right\} \end{bmatrix}. \tag{3.2}$$

Inspection of Equation 3.2 suggests a possible difficulty during analysis given the linear relationship between the equations for δ and P_M. This difficulty arises from inclusion of an isochronous governor in this relatively simple two-bus network. The difficulty is obviated by setting $P_M = \left(-G/2\pi 60\right)\delta$, which has the effect of removing the equation for P_M. This results in the following revised equation for the state variables:

$$\dot{x} = \begin{bmatrix} 2\pi 60\,\Delta\omega \\ \frac{1}{2H}\left\{-\frac{G}{2\pi 60}\delta - P_G - D\Delta\omega\right\} \\ \frac{1}{T'_{d0}X'_d}\left\{-X_d E'_q + X'_d E_f + \left(X_d - X'_d\right)V_1\cos(\delta-\theta_1)\right\} \\ \frac{1}{T_L}\left\{-P_L + (1-K_{LS}\,q_{LS})\,P_0(t)\left(\frac{V_2}{V_{ref}}\right)^{\alpha}\right\} \\ \frac{1}{T_L}\left\{-Q_L + (1-K_{LS}\,q_{LS})\,Q_0(t)\left(\frac{V_2}{V_{ref}}\right)^{\beta}\right\} \end{bmatrix}, \tag{3.3}$$

where $\dot{x}$ is now reduced by one variable:

$$\dot{x} = \begin{bmatrix} \dot{\delta} \\ \Delta\dot{\omega} \\ \dot{E}'_q \\ \dot{P}_L \\ \dot{Q}_L \end{bmatrix}. \tag{3.4}$$

The constraint equations are

$$0 = \begin{bmatrix} G_{12}V_1^2 - V_1V_2\left(G_{12}\cos\theta_{12} + B_{12}\sin\theta_{12}\right) - P_G \\ -B_{12}V_1^2 - V_1V_2\left(G_{12}\sin\theta_{12} - B_{12}\cos\theta_{12}\right) - Q_G \\ -G_{12}V_2^2 + V_1V_2\left(G_{12}\cos\theta_{21} + B_{12}\sin\theta_{21}\right) - P_L \\ B_{12}V_2^2 + V_1V_2\left(G_{12}\sin\theta_{21} - B_{12}\cos\theta_{21}\right) - Q_L \end{bmatrix}. \tag{3.5}$$

As before,

$$\begin{aligned} P_G &= \frac{E'_qV_1}{X'_d}\sin\left(\delta-\theta_1\right) + \frac{V_1^2}{2}\left(\frac{1}{X_q} - \frac{1}{X'_d}\right)\sin\left(2\delta - 2\theta_1\right), \text{ and} \\ Q_G &= \frac{E'_qV_1}{X'_d}\cos\left(\delta-\theta_1\right) - V_1^2\left(\frac{\sin^2\left(\delta-\theta_1\right)}{X_q} + \frac{\cos^2\left(\delta-\theta_1\right)}{X'_d}\right). \end{aligned}$$

We have written for convenience

$$\begin{aligned} \theta_{12} &= \theta_1 - \theta_2\,, \\ \theta_{21} &= \theta_2 - \theta_1\,, \\ G_{12} &= q_{LT1}G_1 + q_{LT2}G_2\,, \text{ and} \\ B_{12} &= q_{LT1}B_1 + q_{LT2}B_2\,. \end{aligned}$$

The system outputs are

$$y = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \delta \\ \Delta\omega \\ E'_q \\ P_L \\ Q_L \end{bmatrix}. \tag{3.6}$$

The simulation parameters used are as follows (in per-unit unless specified). The generator electrical parameters are

$$X_d = 0.9\,, \qquad X_q = 0.8\,, \qquad X'_d = 0.3\,, \qquad T'_{d0} = 7\,\text{s}\,, \qquad E_f^{min} = 0\,, \text{ and } E_f^{max} = 5\,.$$

The generator mechanical parameters are

$$\omega_S = 2\pi 60\,\text{rad/s}\,, \qquad H = 3\,\text{s}\,, \qquad D = 0.3\,, \text{ and } G = 1\,.$$

The line parameters are

$$R_1 = 0.1\,, \qquad X_1 = 0.2\,, \qquad R_2 = 0.05\,, \text{ and } X_2 = 0.1\,.$$

The load parameters are

$$V_{ref} = 1\,, \qquad T_L = 60\,\text{s}\,, \qquad \alpha = 1.5\,, \qquad \beta = 2.5\,, \text{ and } K_{LS} = 0.05\,.$$

E_f for the generator is set to 1.01. Using these values, the initial conditions and values for the system variables at the stable equilibrium point used in the analysis are

$$x = \begin{bmatrix} -70.3477 \\ 0 \\ 0.9685 \\ 0.1853 \\ 0.0352 \end{bmatrix}, \; z = \begin{bmatrix} 0.9594 \\ 0.9503 \\ -70.5034 \\ -70.5157 \end{bmatrix}, \; u = \begin{bmatrix} 1.01 \end{bmatrix}, \; u_d = \begin{bmatrix} 0.2 \\ 0.04 \end{bmatrix}, \text{ and } q = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}.$$

The simplification for P_M (here, equal to 0.1866) means that δ lies well outside the normal range of 0 to 2π. Of course, the critical aspect of the angles are their differences, which remain reasonable.

3.2 Test System Analysis

The system has been constructed such that it has acceptable operation given the stated low demand conditions only when both lines are in service. If the acceptable region for the variables is represented as

$$V_1, V_2 \;\geq\; 0.94\,\text{pu}, \tag{3.7}$$

then the analysis question becomes: for what u_d or q are grid conditions acceptable (excepting that q_{LS} is held at zero)? In terms of cyber security, we ask: can a cyber attacker cause an undesirable impact by successfully tripping either transmission line? The first step of this analysis is construction of the FSA model for the two-bus system, which is the work accomplished during the FY08 period.

Here we exploit recent developments in control systems theory [20] that permit construction of the FSA (a kind of discrete-state machine) that contains the transitions between the discretized states of the dynamical system. The construction guarantees that if the dynamical system takes a

neighborhood of one state into the neighborhood of another, the FSA does the same thing, and vice versa. The discretization employed here is a simple hypercube in the space of dynamical variables. For linear systems a fast code has been developed (Algorithm 3.2 in [20]) that produces an FSA that contains the transitions between the states.

The analysis requires that we translate the differential-algebraic equations (i.e., constrained dynamics) into an unconstrained ordinary differential equation. We employ singular perturbation theory [21] to formally relax the constraints, i.e., we introduce pseudo-velocities for the former constraint variables and take the limit as they go to zero, as follows:

$$M \begin{bmatrix} \dot{\delta} \\ \Delta\dot{\omega} \\ \dot{E}'_q \\ \dot{P}_L \\ \dot{Q}_L \\ \dot{V}_1 \\ \dot{V}_2 \\ \dot{\theta}_1 \\ \dot{\theta}_2 \end{bmatrix} = \begin{bmatrix} 2\pi 60\,\Delta\omega \\ \frac{1}{6}\left\{\frac{\delta}{2\pi 60} - P_G - 0.3\,\Delta\omega\right\} \\ \frac{1}{2.1}\left\{-0.9\,E'_q + 0.303 + 0.6\,V_1\cos(\delta-\theta_1)\right\} \\ \frac{1}{60}\left\{-P_L + P_0(t)\,V_2^{1.5}\right\} \\ \frac{1}{60}\left\{-Q_L + Q_0(t)\,V_2^{2.5}\right\} \\ G_{12}V_1^2 - V_1V_2\left(G_{12}\cos\theta_{12} + B_{12}\sin\theta_{12}\right) - P_G \\ -B_{12}V_1^2 - V_1V_2\left(G_{12}\sin\theta_{12} - B_{12}\cos\theta_{12}\right) - Q_G \\ -G_{12}V_2^2 + V_1V_2\left(G_{12}\cos\theta_{21} + B_{12}\sin\theta_{21}\right) - P_L \\ B_{12}V_2^2 + V_1V_2\left(G_{12}\sin\theta_{21} - B_{12}\cos\theta_{21}\right) - Q_L \end{bmatrix} \tag{3.8}$$

where

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \varepsilon_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon_3 \\ 0 & 0 & 0 & 0 & 0 & 0 & \varepsilon_4 & 0 & 0 \end{bmatrix} \tag{3.9}$$

and $|\varepsilon_j| \to 0$. The assignment of both the signs of the ε_j and the new state variables to the extended vector field is important; we rely on inspection for the former[1] and the discussion in [22] and [23] for the latter. In the new formulation,

$$X = \begin{bmatrix} x \\ z \end{bmatrix} \text{ and } U = u_d\ .$$

[1] We chose $\varepsilon_1, \varepsilon_3 < 0$ and $\varepsilon_2, \varepsilon_4 > 0$ to guarantee a stable system at equilibrium.

Note that $u = E_f = 1.01$ has been fixed, and the states for q are addressed by enumerating different dynamical systems for Equation 3.8. To implement the deterministic finite automaton (DFA) analysis, we linearized Equation 3.8 around the equilibrium as follows:

$$\dot{X}(t) = A\,\Delta X + B\,\Delta U \tag{3.10}$$

where

$$\Delta X = X(t) - X_{equilibrium},$$

$$\Delta U = U(t) - U_{equilibrium},$$

A is the Jacobian of the right-hand-side of Equation 3.8 with respect to X evaluated at $X_{equilibrium}$, and B is the Jacobian with respect to U evaluated at $U_{equilibrium}$. Such linearization is routine for stability analysis and is not expected to affect our conclusions about the vulnerability of the system.

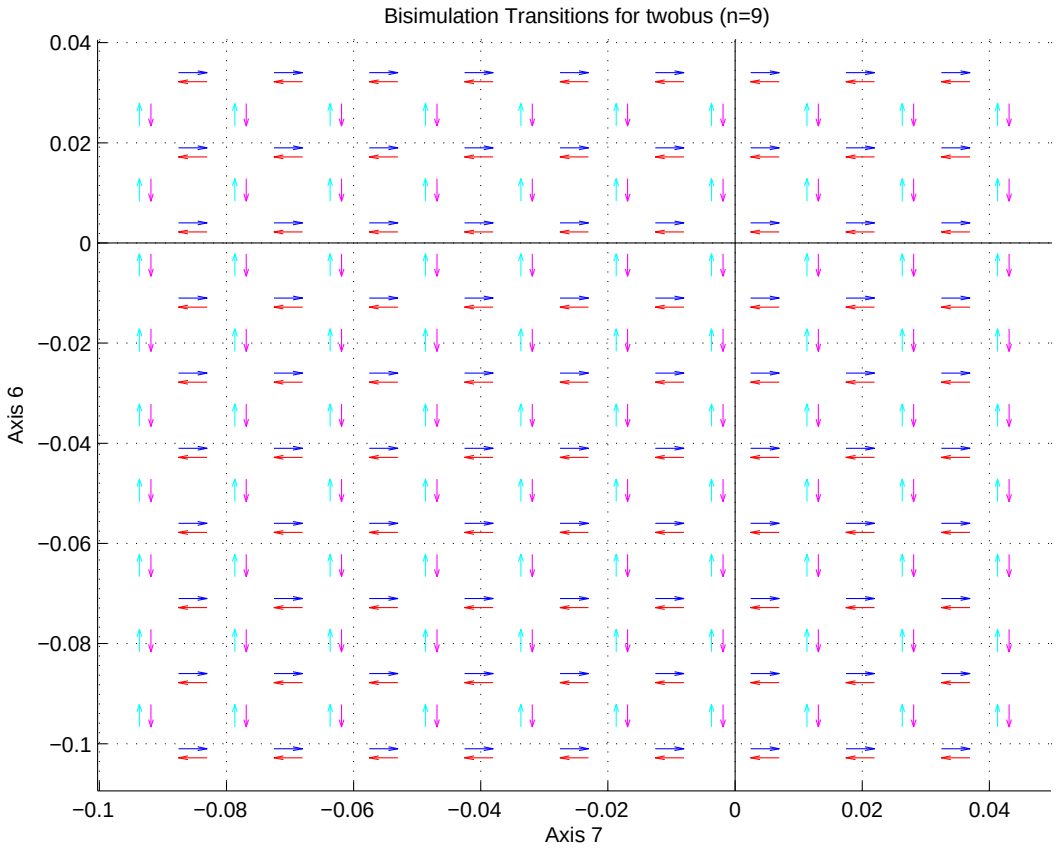

Figure 3.2: Bisimulation transitions for V_1 (Axis 6) and V_2 (Axis 7).

Vulnerability analysis proceeds from the FSA as follows. If one can identify a set of nominal states and a set of failure states (the key difference being the criterion in Equation 3.7), one can ask if there exists a directed path in the FSA that connects the set of nominal states with the set of failed states. If no such path exists, there are no inputs that will lead to failure from nominal conditions. On the other hand, the existence of the path means that there exist some inputs that will take a nominal state to failure (under varying conditions for q); for this model failure consists of the plane defined by $V_1 \leq 0.94$ and $V_2 \leq 0.94$. In this model such paths exist (as expected) when one or the other line is tripped; but we also found that even with no lines tripped there exist some inputs that take an element of the set of nominal states somewhere into the set of failed states. The situation is depicted in Figure 3.2, where $q_{LT1} = q_{LT2} = 1$ (i.e., both lines are in service), Axis 6 corresponds to V_1, Axis 7 corresponds to V_2, and the equilibrium state is at the origin. The figure represents a subsection of the nine-dimensional hypercube formed by the FSA analysis performed on Equation 3.8.

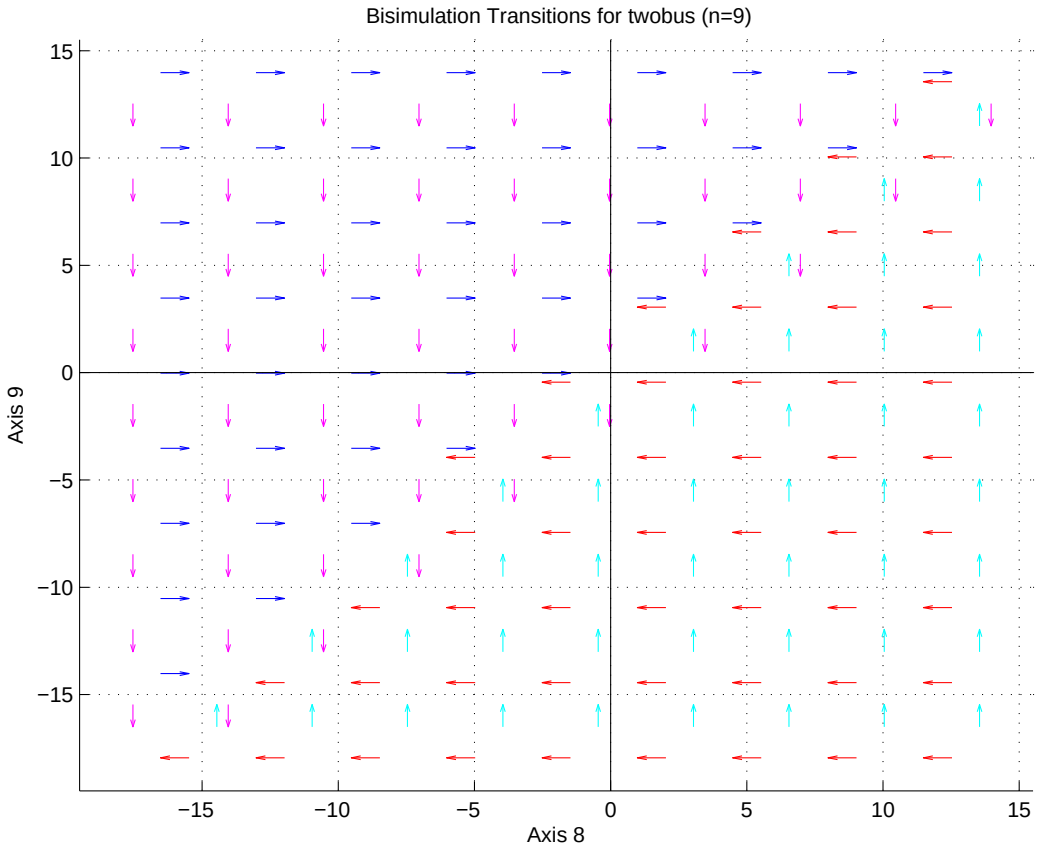

Figure 3.3: Bisimulation transitions for θ_1 (Axis 8) and θ_2 (Axis 9).

Our analysis for this phenomenon (reachability of the unacceptable region with both lines in service) is that the current analysis capability for the FSA algorithm allows bounding on the input variables (here, u_d only, as we fixed $u = E_f = 1.01$ to eliminate it from consideration). Given this stipulation, it is obvious that an unbounded u_d makes most voltage states reachable, and all of them in fact over the extent of Figure 3.2. Part of the planned work for FY09 includes applying a bound for u_d, so that its range is no longer arbitrary, and reachability may be calculated under more reasonable conditions.

For reference we also show a non-trivial transition matrix for the θ_1 and θ_2 plane (Figure 3.3). Although in this plane there are no states corresponding to failure, the bisimulation correctly displays the tendency of the dynamics to drive the angles toward equality (in this system, the relatively low impedance of the two lines between buses one and two results in roughly equal phase angles for their voltages).

3.3 Section Conclusions

Finite state abstraction (FSA) of the continuous dynamics of a power model offers a rigorous assessment of the vulnerability of the modeled system while at the same time remaining computationally tractable. Two items remain for future work, aspects of which are discussed in the review by Colbaugh et al. [20]:

- The results obtained so far assess the vulnerability for any and all possible inputs. The utility of such results could be increased by restricting the range of possible inputs. This would in turn require a reformulation of the FSA. Such a reformulation seems achievable to us but would require that we compute a simulation rather than a bisimulation, the difference being that failed states reachable in the FSA would not necessarily be reachable in the original system; any vulnerability identified in the FSA would need to be individually rechecked in the original system. On the other hand, states unreachable in the FSA would continue to be unreachable in the original system and would not require rechecking [24].
- So far we need to recalculate and check the FSA for each combination of lines closed and open. For this model we need to check only three configurations (both open, one open and two closed, one closed and two open), but the number of configurations to be checked grows rapidly with the number of lines. Therefore, the algorithm ought to incorporate switching between configurations into the analysis at the beginning. This also requires a reformulation of the FSA that we think is feasible.

Chapter 4

Report Conclusions

The two approaches taken in this research, reliability impacts from cyber attack RICA and finite state abstraction FSA, provide a foundation for the *quantitative* evaluation of impacts to the power grid caused by cyber attacks. The RICA method allows calculation of the degradation in reliability caused by cyber attacks. The analysis can be understood as a means to calculate *averaged* measures for the ongoing value of good cyber security (or alternatively, the cost of deficient security). It depends on effective and agreed-upon characterization for the adversary and the effects caused by cyber attacks. This calculation allows utility managers to estimate the cost of unmitigated vulnerabilities, and to plan budgeting for their remediation.

The FSA research complements the RICA algorithm, as it provides a potential path to solving the previously intractable problem of determining which cyber attacks can cause significant problems. The algorithm has been improved to the point where the finite-state model can be calculated for simple systems, such that simple cyber vulnerability analysis is beginning to become feasible. Future work will develop organized approached for evaluating vulnerabilities and risks.

Chapter 5

Recommendations

To improve on the Impacts Analysis work, we recommend the following additional research. These are categorized pertaining to the research approach.

For Reliability Impacts from Cyber Attack RICA:

1. Expand the testing to include larger power system models.
2. Investigate the parametric sensitivities for the restoration times.
3. Determine methods to improve the estimates for mean time to attack MTTA.
4. Research opportunities to include more advanced (though still probabilistic) attack models.

For Finite State Abstraction FSA:

1. Constrain the inputs to be bounded.
2. Develop a method to integrate the system switching behavior into a single model.
3. Determine an algorithm to extract information for the paths to unacceptable operation from the FSA.
4. Perform tests on a larger system.

Appendix A

References

[1] J. Stamp, R. Colbaugh, R. Laviolette, A. McIntyre, and B. Richardson, *Final Report: Impacts Analysis for Cyber Attack on Electric Power Systems (National SCADA Test Bed FY07)*, Sandia National Laboratories report SAND2008-7066P (July 2008).

[2] M. Schilling, A. L. De Silva, R. Billinton, and M.A. El-Kady, "Bibliography on Power System Probabilistic Analysis (1962-1988)," in *IEEE Transactions on Power Systems*, Vol. 5, No. 1, pp. 41-49 (February 1990).

[3] R. Allan, R. Billinton, A. Breipohl, and C. Grigg, "Bibliography on the Application of Probability Methods in Power System Reliability Evaluation 1967-1991," in *IEEE Transactions on Power Systems*, Vol. 9, No. 1, pp. 41-49 (February 1994).

[4] R. Allan and R. Billinton, "Reliability Assessment of Composite Generation and Transmission Systems," IEEE Power Engineering Society Tutorial, 90EH0311-1-PWR (1989).

[5] J.R. Ubeda and R. Allan, "Sequential Simulation Applied to Composite System Reliability Evaluation," in *IEE Proceedings C*, Vol. 139, No. 2, pp. 81-86 (March 1992).

[6] R. Allan and R. Billinton, "Probabilistic Assessment of Power Systems," in *Proceedings of the IEEE*, Vol. 88, No. 2 (February 2000).

[7] B. Littlewood, S. Brocklehurst, N. Fenton, P. Mellor, S. Page, D. Wright, J. Dobson, J. McDermid, and D. Gollmann, "Towards Operational Measures of Computer Security," in *Journal of Computer Security*, Vol. 2, No. 3, pp. 211-229 (1993).

[8] K. J. Soo Hoo, "How Much Is Enough? A Risk-Management Approach to Computer Security," technical report, Consortium for Research on Information Security and Policy (CRISP), Stanford University, (June 2000).

[9] B. B. Madan, K. Goševa-Popstojanova, and K. Vaidyanathan, "Modeling and Quantification of Security Attributes of Software Systems," in *Proceedings of the International Conference on Dependable Systems and Networks*, pp. 505-514, (June 2002).

[10] C. Taylor, A. Krings, and J. Alves-Foss, "Risk Analysis and Probabilistic Survivability Assessment (RAPSA): An Assessment Approach for Power Substation Hardening," in *Proceedings of the 1st Workshop on Scientific Aspects of Cyber Terrorism*, (November 2002).

[11] S. Singh, M. Cukierz, and W. H. Sanders, "Probabilistic Validation of an Intrusion-Tolerant Replication System," in *Proceedings of the International Conference on Dependable Systems and Networks*, pp. 615-624, (June 2003).

[12] J. R. Conrad, "Analyzing the Risks of Information Security Investments with Monte-Carlo Simulations," in *Proceedings of the 4th Workshop on the Economics of Information Security*, Kennedy School of Government, Harvard University (June 2005).

[13] E. Jonsson and T. Olovsson, "A Quantitative Model of the Security Intrusion Process Based on Attacker Behavior," in *IEEE Transactions on Software Engineering*, Vol. 23 No. 4, p. 235 (April 1997).

[14] J. McDermott, A. Kim, and J. Froscher, "Merging Paradigms of Survivability and Security: Stochastic Faults and Designed Faults," in *Proceedings of the Workshop on New Security Paradigms*, pp. 19-25, ACM:New York (2003).

[15] N. F. Schneidewind, "Reliability – Security Model," in *Proceedings of the 11th International IEEE Conference on Engineering of Complex Computer Systems (ICECCS'06)*, pp. 269-278 (August 2006).

[16] J. McDermott, "Attack-potential-based Survivability Modeling for High-consequence Systems," in *Proceedings of the 3rd International Workshop on Information Assurance*, pp. 119-130 (March 2005).

[17] S. Staniford, V. Paxson, and N. Weaver, "How to 0wn the Internet in Your Spare Time," in *Proceedings of the 11th Usenix Security Symposium*, pp. 149-167 (2002).

[18] R. Billinton, S. Kumar, N. Chowdhury, K. Chu, K. Debnath, L. Goel, E. Khan, P. Kos, G. Nourbakhsh, J. Oteng-Adjei, "A Reliability Test System for Educational Purposes – Basic Data," in *IEEE Transactions on Power Systems*, Vol. 4, No. 3, pp. 1238-1244 (August 1989).

[19] R. Billinton and H. Yang; "Incorporating Maintenance Outage Effects in Substation and Switching Station Reliability Studies," in *Proceedings of the Canadian Conference on Electrical and Computer Engineering*, pp. 599-602 (May 2005).

[20] R. Colbaugh, K. Glass, and G. Willard, *Analysis of Complex Networks Using Aggressive Abstraction*, Sandia National Laboratories report SAND2008-7327 (October 2008).

[21] P. V. Kokotovic, H. K. Khalil, and J. O'Reilly, *Singular Perturbation Methods in Control: Analysis and Design*, SIAM:Philadelphia (1986).

[22] E. Barany, S. Schaffer, K. Wedeward, and S. Ball, "Nonlinear Controllability of Singularly Perturbed Models of Power Flow Networks," in *Proc. 43rd IEEE Conference on Decision and Control*, pp. 4826-4832 (2004).

[23] M. Ilic and J. Zaborsky, *Dynamics and Control of Large Electric Power Systems*, Wiley-Interscience:New York (2000).

[24] N. Giorgetti, G. Pappas, and A. Bemporad, "Bounded Model Checking of Hybrid Dynamical Systems," in *Proceedings of the 44th IEEE European Conference on Decision and Control*, pp. 672-677 (December 2005).

Appendix B

Acronyms, Symbols, and Abbreviations

Table B.1: Acronyms

Acronym	Phrase
APTB	average percentage of tripped breakers
C2P	cyber-to-physical
DOE	Department of Energy
DOI	duration of interruption
EIR	energy index of reliability
ENSI	energy not served per interruption
EPG	electric power grid
FOI	frequency of interruption
FSA	finite state abstraction
FY	fiscal year
IA	impacts analysis
LCI	load curtailed per interruption
LOEE	loss of energy expectation
LOLE	Loss of Load Expectancy
LLD	load loss duration
MC	Monte-Carlo
MTTA	mean time to attack
MTTB	mean time to breach
MTTF	mean time to failure
MTTR	mean time to recover
NERC	North American Electric Reliability Corporation
NSTB	National SCADA Test Bed
OPF	optimal power flow
PEPA	performance evaluation process algebra
PRA	probability risk assessment
RBTS	Roy Billinton test system
RICA	reliability impacts from cyber attack
ROCOF	rate of occurrence of failures
SCADA	supervisory control and data acquisition

Table B.2: Symbols

Symbol	Units	Description
α	none	load voltage exponential variability for active power
A	varies	Jacobian of the system with respect to X at equilibrium
β	none	load voltage exponential variability for reactive power
B	varies	Jacobian of the system with respect to U at equilibrium
B_{ij}	per-unit	susceptance from bus i to bus k
C	varies	matrix relating state variables to observable outputs
δ	radians	electrical angle between its voltage and the rotor major axis
D	per-unit	generator rotor damping constant
$\Delta\omega$	per-unit	relative speed of the rotor and stator field
ΔX	varies	difference between state variables and equilibrium values
ΔU	varies	difference between control variables and equilibrium values
E_f	per-unit	generator field circuit voltage
E_f^{max}	per-unit	maximum generator field
E_f^{min}	per-unit	minimum generator field
ε_j	none	pseudo-velocity constant for relaxed constraints
E'_q	per-unit	q-axis generator induced voltage
f	varies	system differential equations
g	varies	system algebraic equations
G	per-unit	gain constant of a generator governor
G_{ij}	per-unit	conductance from bus i to bus k
H	per-unit	generator rotor inertia constant
K_{LS}	none	discrete load shed factor
λ_{GEN}	occ/yr	rate of successful generator protection attack
λ_{LINE}	occ/yr	rate of successful line protection attack
λ_{SCADA}	occ/yr	rate of successful SCADA attack
M	none	multiplicative matrix for singular perturbation formulation
N	none	the number of elements in a power grid model
ω_R	per-unit	machine rotor speed
ω_S	per-unit	speed of a machine's stator field
π	none	ratio of the circumference to the diameter of a circle
$P_0(t)$	per-unit	driving signal for load active power demand
P_G	per-unit	generator active electrical power
P_L	per-unit	load active power
P_M	per-unit	generator mechanical power
p_{SCADA}	none	percentage of breakers tripped per SCADA attack
q	varies	vector of discrete control variables
$Q_0(t)$	per-unit	driving signal for load reactive power demand

Continued on next page

Table B.2 – *continued from previous page*

Symbol	Units	Description
Q_G	per-unit	generator reactive power
Q_L	per-unit	load reactive power
q_{LS}	none	discrete current load shed level variable
q_{LTi}	none	discrete trip variable for line i
R_k	per-unit	resistance for grid element k
S_G	per-unit	generator complex power
S_L	per-unit	load complex power
θ_i	radians	phasor voltage phase angle of bus i
θ_{ij}	radians	phasor voltage phase angle between buses i and j
T'_{D0}	seconds	open circuit time constant for the generator d-axis
T_L	seconds	load recovery time constant (in per unit)
u	varies	vector of continuous control variables
U	varies	vector of control variables after singular perturbation translation
u_d	varies	vector of system disturbance functions
$U_{equilibrium}$	varies	vector of control variables at equilibrium in singular perturbation form
V	per-unit	phasor voltage magnitude
V_{ref}	per-unit	reference voltage for load exponential variability
x	varies	vector of state variables
$\dot{x}$	varies	first derivative with respect to time of state variable x
X	varies	vector of state variables after singular perturbation translation
$\dot{X}$	varies	first derivative with respect to time of state variable X
X_d	per-unit	d-axis reactance of a generator
X'_d	per-unit	d-axis transient reactance of a generator
$X_{equilibrium}$	varies	vector of unconstrained state variables at equilibrium
X_k	per-unit	reactance for grid element k
X_q	per-unit	q-axis reactance of a generator
y	varies	vector of system outputs
z	varies	vector of algebraic variables

Table B.3: Abbreviations

hr	hour(s)
int	interruption
MW	megawatts
occ	occurrences
pu	per-unit
rad	radians
s	seconds
yr	year(s)

Appendix C

Glossary

Table C.1: Definitions

Term	Definition
reliability impacts from cyber attack (RICA) algorithm	The RICA approach presented and developed in this report is intended to extend the conventional study of power grid reliability to include the impacts of cyber attack against control systems. In particular, given an agreed-upon set of assumptions regarding the likely rate of successful control system attacks and their expected recovery periods, the RICA approach can calculate the expected reduction in system reliability caused by cyber attack; this can allow a value to be placed on the modeled level of cyber insecurity.
finite state abstraction (FSA) algorithm	The FSA approach represents a possible method to determine the stability and suitability of operations for an electrical power grid utilizing a fully descriptive model (including dynamics and control action) under a wide range of potential cyber attack scenarios. The underlying premise for this work is the conversion of the hybrid grid/control model dynamics into a representative FSA, which greatly improves the tractability of the problem. The new approach is expected to scale to larger systems and modeling complexity much better than traditional analysis for grid dynamics.

Appendix D

Contacts

Table D.1: For More Information

Name	Organization
Jason Stamp *Sandia Project Lead*	Sandia National Laboratories P.O. Box 5800 Albuquerque, NM 87185-1108 jestamp@sandia.gov
Bob Pollock *Sandia NSTB Lead*	Sandia National Laboratories P.O. Box 5800 Albuquerque, NM 87185-0671 rdpollo@sandia.gov
Jennifer Depoy *Sandia NSTB Manager*	Sandia National Laboratories P.O. Box 5800 Albuquerque, NM 87185-0671 jmdepoy@sandia.gov
Hank Kenchington *DOE NSTB Manager*	U.S. Department of Energy 1000 Independence Avenue, SW Washington, DC 20585 henry.kenchington@hq.doe.gov

Appendix E

Distribution

Table E.1: Distribution

Name	Location
Jason Stamp *Sandia Project Lead*	Sandia National Laboratories MS 1108 Albuquerque, NM 87185-1108
Bob Pollock *Sandia NSTB Lead*	Sandia National Laboratories MS 0671 Albuquerque, NM 87185-0671
Jennifer Depoy *Sandia NSTB Manager*	Sandia National Laboratories MS 0671 Albuquerque, NM 87185-0671
Hank Kenchington *DOE NSTB Manager*	U.S. Department of Energy 1000 Independence Avenue, SW Washington, DC 20585
Sandia National Laboratories Technical Library	Sandia National Laboratories MS 0899 Albuquerque, NM

Sandia National Laboratories