A SENSITIVITY ANALYSIS OF CYBER CONTINGENCY RANKING WITHIN THE SOCCA FRAMEWORK

BY

AARON ROBERT PHELPS

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2014

Urbana, Illinois

Adviser:

Research Assistant Professor Rakesh Bobba

ABSTRACT

Cyber-infrastructure is at the heart of power system operations and is critical for maintaining reliable and stable power supply. The advent of smart grid technology will undoubtedly increase the exposure and potential avenues of cyber attack well into the future. The industry standard of contingency analysis largely focuses on accidental outages, such as natural disasters, equipment malfunction, etc. Intentional, directed attacks and cyber components are not well understood or accounted for.

In response, the Security-Oriented Cyber-Physical Contingency Analysis (SOCCA) framework demonstrates that it is both prudent and practical to assess the impact of cyber events within power infrastructures. Using a new formalism to model cyber-physical interconnections and by ranking contingencies based on impact and attack complexity, SOCCA presents system operators with a detailed vulnerability landscape of their networks. SOCCA's contingency ranking algorithm relies heavily on Markov Decision Processes. These MDPs require expert knowledge in determining the attack surface and gauging the likelihood of an attack's success as represented by a probability. The choice of reward function and

assignment of probabilities greatly influence the behavior of the MDP. Therefore, the accuracy of the ranking algorithm is called into question as it is intrinsically tied to the accuracy of the expert knowledge.

This thesis aims to identify the major factors that affect the contingency ranking in an MDP model that represents an industry-standard cyber-physical power network. Probability assignments will be varied including augmenting the SOCCA framework in order to extend the probabilities associated with the MDPs to be bounded intervals rather than exact values. This way, reliance on precise expert knowledge is lessened and sensitivity analysis can be performed to provide a confidence rating to the contingency analysis. This will also give insight into how modifying or mitigating certain attack steps contributes to the overall cyber security state of the network.

ACKNOWLEDGEMENTS

The author wishes to extend his utmost gratitude to Professor Rakesh Bobba for his continuous support and sagacious advice over the last 18 months. It has truly been a pleasure and a privilege to work alongside you. Special thanks to Professor Saman Zonouz and Luis Garcia from the University of Miami for their technical expertise and assistance in code development. Thanks also to Kate & Matt Davis from the PowerWorld Corporation for providing key knowledge of power simulation and power networks. Last but not least, many thanks to Dr. Masooda Bashir, Syed Faisal Hasan and the rest of the ICSSP family for not only funding to complete my graduate degree, but for the encouragement, excitement, and inspiration along the way.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION

Called mankind's biggest and most complex machine, the modern electric power grid is a myriad of aging technologies and ad-hoc solutions. Assessing the cyber threat landscape of this most critical asset is vital in ensuring its reliability and the continued delivery of the nation's essential services. Traditional state estimation in the power grid refers to analyzing data from a distributed network of sensors connected to power components and communicated through SCADA or similar technologies [1]. These sensors operate on line voltages, current, phase angle and discrete measurable electricity phenomena.

This implies that the behavior of the system can be accurately predicted using power flow equations, using numerical methods largely based on the Newton-Raphson method [2]. This allows power companies to plan for and mitigate the damage caused by disruptive events during operation and know the state of the system of all times. This usually manifests itself as 'N – 1' contingency analysis [3], where the loss of a generator or transmission component is simulated and system operators make

adjustments to re-route power as necessary. NERC standards require such capability.

However, as smart grid technology continually pushes its way into the electric sector, the industry can no longer rely on purely physical numerical methods to guarantee reliability. Promising big improvements in energy efficiency and benefits to utilities and consumers, computer automation of the grid also ensures increased exposure to the risk of outages from cyber threats. Internet-ready smart meters under consumer control, distributed generation and energy storage, and GPS-enabled phasor measurement units are just a few examples of how the cyber attack landscape may grow unabated for many years to come.

1.1 The SOCCA Framework

In order to address the need to incorporate cyber-physical interconnections into power grid contingency analysis, Zonouz et al. presented SOCCA: A Security-Oriented Cyber-Physical Contingency Analysis in Power Infrastructures [4]. First, SOCCA uses offline discovery to map the cyber control network. An XML document was created to describe the topology and network access policies / firewall rules. Using the NetAPT tool [5], a connectivity

matrix is created using the topology and access policies that reflects possible workflows through the network as a directed graph. In the future, this could be done in a dynamic, online manner using tools such as nmap, Nessus, etc.

The power component network is handled by the Power World Simulator [6], which performs power flow contingency analysis to determine component failures that would result in an outage. The interconnections between the cyber control network and the power component network are then added. This is typically a breaker connected to a relay, which is handled by a front end processer in the control network. Using the complete directed connectivity graph, SOCCA generates a Markov Decision Process [7] model that seeks to enumerate the paths an attacker might take in securing resources and privileges in an attempt to cause line outages.

For each set of privileges that an attacker may possess, a unique state is given in the MDP. Using the Power World simulation, each state is granted a performance index, which signifies the outage severity associated with disruption of the network components governed by those privileges. Next, using value iteration, a security index for each state is calculated using the performance

index (impact) with the probability that the attacker will succeed in gaining the state's privileges from cyber attack. These probabilities are, at best, speculative and based on expert knowledge or experience. Adding mitigation schemes and network defenses likely serve to lower the likelihood of success, but it is hard to evaluate by how much this will change. An attacker may also discover a novel approach to infiltration that renders defenses ineffective or take advantage of social engineering and phishing to achieve his/her goals. It becomes crucial when assessing critical assets and contingencies to know how sensitive the model is to likelihood metrics.

Once the value iteration is complete, the attacker's possible actions are ranked according to the optimal route the attacker can take to cause disruption. Because the MDP value iteration algorithm incorporates all future actions into the analysis, ranking is accomplished by simply choosing the action with highest security index from the attacker's available actions given the set of privileges he/she has attained. It is important to note that as new information comes in regarding cyber events as reported by intrusion detection systems or other network monitoring, SOCCA can be run in real-time to re-evaluate from any point in the MDP

using the attacker's current position as the base point. This is helpful as the network grows and scales and performance constraints limit the ability to generate and evaluate the entire MDP at once. Instead, the contingency analysis will be limited to a given horizon.

Using the contingency rankings generated by the MDP, a system operator has good knowledge of the most crucial assets on his system and which intrusions are the most important to prevent against at varying stages of network operation. Tools that work in real-time and report early are vital due to timing deadlines inherent to power infrastructure. System operators typically know their options to re-route power in the event of a failure, and they can use the same tactics in the event of a cyber attack.

1.2 Markov Decision Processes

It is important to understand the manner in which MDPs model decision-making scenarios to give the results of the contingency analysis clear meaning. A Markov Decision Process [7] is a tuple $(S, A, P_a(s,s'), R(s))$ where S is a finite set of states, A is a finite set of actions, P is a probability transition function that reflects the likelihood that given an action a taken from state s, the system

changes to state *s'*, and *R* is the expected reward gained from for achieving state *s*. Since *S* and *A* are finite, an MDP of this type is referred to as a discrete Markov decision process.

For the purposes of this thesis, the MDP is modeled after the SOCCA framework [4]. The state space *S* is the security state of the cyber-power network, meaning that each state is a potential set of privileges that an attacker has gained through infiltration. A privilege in this sense is defined as control over a network host or power component. For each new node in the network that an attacker controls, he/she is said to have extended the current set of privileges and the MDP transitions to a new state. The attacker initially has no privileges in the network and is seen as a user on the open Internet.

The action space *A* represents the attacker's choices as he/she explores the security state space of the system. Given the current set of privileges, any directed connectivity to another node that an attacker could potentially exploit is assigned an action that transitions to a new state of privileges including the exploited node. Taking an action can result in two possibilities. An attacker is either successful at exploiting the node and gains the new

privilege with probability *p* or fails to execute the attack and remains at his/her current state with probability (1-*p*). These probabilities are what define *P*. An example MDP is shown in Figure 1 [4].

Lastly, *R*, the reward function, reflects how transitioning from states in the MDP serves the attacker's goals. In this case, the attacker's goal is characterized as causing as many line outages as he/she can. For each state *s,*

$$R(s) = \sum_{l \epsilon L} \left[\max\{ \frac{f_s(l)}{f^{MAX}(L)} - 1, 0 \} \right]^2$$

which says that for every power line *l*, calculate the flow on that line divided by the maximum flow permitted on that line and if there is a line flow violation, assign a positive value. Otherwise, assign 0. If the security state of the system *s* contains privileges for components that directly control flow on power lines, it is assumed that the attacker compromises those lines and switches them off. The flow of power under each set of outages is calculated within the Power World Simulator.

Once the MDP has been properly filled, it is ready to find the optimal policy for the decision maker. A policy, denoted $\prod(s)$, is

the choice of action that the decision maker takes while at state *s*. The decision maker here is the attacker. To achieve this, the dynamic programming algorithm called value iteration is used. For value iteration, a value function V(s) is defined that gives each state *s* a value *v* so that it can be compared amongst other states. This differs from the reward function in that it mixes the reward gained at a state with the probability of reaching that state as well as the future rewards of states available after reaching *s*. *V(s)* is defined as follows:

$$v(s) = \max_{a \epsilon A} \left\{ F(s) + \gamma \sum_{state\ s'} P_a(s, s') v(s') \right\}$$

Here gamma represents the discount factor, which accounts for the diminished return of rewards that are far away from the current state. As a dynamic programming algorithm, *V(s)* is recalculated for every state iteratively until the sum of change across all states falls within some low value. The value used in this thesis was 10^{-3}. Once values have been calculated for all states, the optimal policy for the attacker to take is the action that would lead to the state of highest value from the current state. Additionally, for the system operator, these actions are the most important to protect and prevent from occurring.

1.3 Overview

From here, Chapter 2 will review the body of research related to cyber contingency analysis in power networks and how it compares to the methods discussed here. Chapter 3 will examine the creation of the cyber-physical power network model and justification for why it is an appropriate representation of modern power infrastructure. It will also dive into the parameters that were assigned to the corresponding MDP. Chapter 4 discusses the methods used to evaluate the dominating factors that affect the contingency ranking and the results of employing those methods. Chapter 5 will draw conclusions about the findings from Chapter 4 and make considerations for analyses to be done in the future.

1.4 Figures

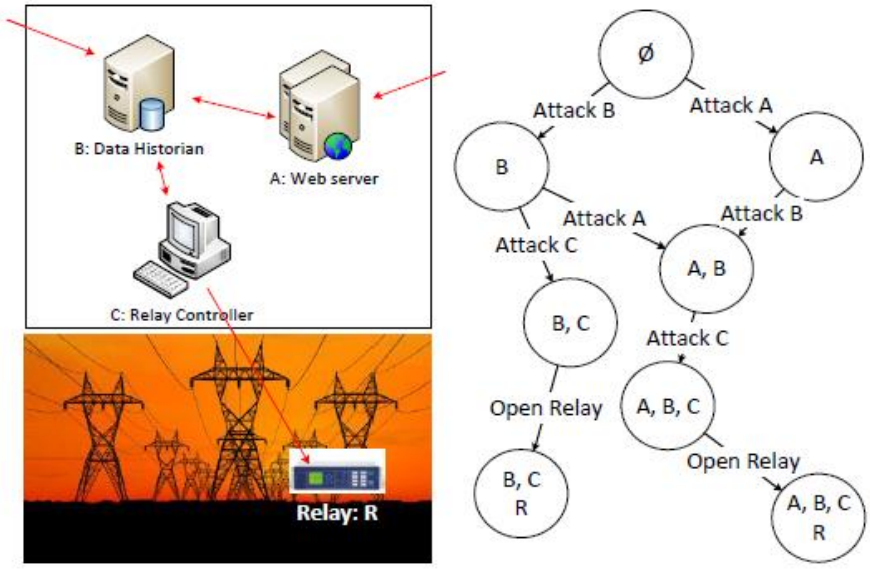

Figure 1: An Example Power Control Network and MDP

CHAPTER 2: RELATED WORK

Although contingency analysis of the power grid has been studied for many decades, only recently have concerted efforts been carried out to incorporate cyber contingencies and their effects into the mix. Chen et al. [8] propose a stochastic Petri net approach to model coordinated cyber-physical attacks on the smart grid. Using a hierarchical method of constructing small Petri nets from specific domain experts, a large Petri net is compiled to represent the system as a whole. The authors admit that there is a great deal of manual effort and expertise required for an accurate model. Additionally, their method focuses on modeling specific attack scenarios such as a smart meter compromise. By shifting focus to the effects of a generic compromise and gauging the impact of the expert knowledge, this research hopes to overcome these limitations.

Research has also been performed that seeks to advance dynamic detection of network compromise and measurement corruption in a power network [9]. While detection is important in its own right, and the authors make good effort to account for unexpected or unknown attacks, their method is complex and intrinsically tied to

a power system state that in reality, morphs and grows constantly. Reevaluation of contingencies in real-time is incredibly important for system operators and focus should be reallocated towards enumeration of key assets and not so heavily on types of attacks.

Sridhar et al. present a detailed overview of cyber security tasks relevant to the power grid domain. While presented at a high-level, the risk assessment methodology explained within contains many vital activities that collectively serve to professionally coordinate cyber defense within the power infrastructure. The SOCCA framework would be an essential contribution to the infrastructure vulnerability analysis within the broader scope of complete cyber-physical system security. As the paper states, grid security must be continually analyzed with respect to complexity, long lifespans, and novel attack vectors. At this time, no other contingency analysis offers the same real-time, scalable protection benefits.

In terms of perturbation analysis of MDPs, different approaches were considered. [10-12] show examples of methods that have been proposed to handle sensitivity of steady-state performance of Markov processes. The approach deemed most closely related to

the problem description at the time of writing is [13], which among other things, contributes a novel value iteration algorithm that allows transition probabilities within an MDP to be treated as a bounded real interval instead of an exact value. This algorithm will be discussed in greater detail in Chapter 4.

CHAPTER 3: MODEL CREATION

Every network and corresponding system model will have significant differences, even within domain-specific applications. Power networks come in all shapes and sizes, governed by the area of distribution they serve, types of equipment in use, utility-specific policies, and legacy issues. Therefore, when advising utilities on the security state of their networks or providing tools to system operators, it is imperative that the model formalism makes sense and emulates reality. There is no perfect generalization or analogous model for every single network and it can make drawing meaningful conclusions difficult in a research setting. This chapter will highlight how the example model was chosen and make a case for having general value to the industry at large.

3.1 The Physical Power Network

When performing cyber-physical contingency analysis and system monitoring, scope must be sufficiently limited so that the system does not become overburdened. Quick decision making makes all the difference in guarding against intentional attacks. That being said, a power network was chosen to reflect the typical granularity at which monitoring is performed – the substation or a small

collections of substations. The simulated power grid infrastructure chosen was the IEEE 24-bus reliability test system [14]. This test system fully describes generation and transmission network, complete with load constraints, ratings, and reliability data. It was, in fact, created in order to sufficiently describe a broad range of power networks and was created by a committee of industry professionals well-versed in real implementations. A diagram of the 24-bus system is shown in Figure 2. The system consists of 32 generation units and 24 load/generation buses connected by 38 transmission lines along with full descriptions of peak load conditions, generation capability, and energy output. This data was directly entered into a Power World simulation in order to calculate highly accurate line flows.

3.2 The Cyber Control Network

As for the cyber control network, there are no peer-reviewed and standard models for the IT infrastructure necessary to coordinate and monitor a power network. In the culture of cyber security awareness that we live in, it is unlikely that any utility would step forward to expose the inner workings of their network willingly to the general public. In place of a standard representation, an example of a control network was constructed from conversations

with local professionals in the industry and based on a real network that was kept confidential. On their recommendations, the cyber control network was made to fit with the 24-bus power network. T

The cyber control network consists of 59 host systems and firewalls and monitors the buses in the 24-bus model. Each bus is controlled by a single host in the control network. Inspection of the cyber control network reveals 10 key assets that affect the cyber-physical interconnection and can result in line outages. Figure 3 shows the cyber control network topology.

3.3 Creation of the Markov Decision Process

Since there are 10 key assets that an attacker would seek privileged access to in order to achieve his/her goals, states are created in the MDP that reflect collection of those 10 privileges. Since the assets must be achieved somewhat sequentially, there are 146 resultant states instead of all possible permutations of the 10 privileges. As described above, the SOCCA framework uses the connectivity matrix and access policy rules to determine the possible actions from each state. The calculation of the reward function has already been discussed. Figure 4 is a directed graph showing the entire MDP. The top number at each state is the ID of

that state, the middle is the state's performance index or reward value, and the bottom number is the state's security index, the result of value iteration. States depicted white contain cyber privileges only while grey states are those that contain physical privileges as well. A mapping from state ID's to control network hosts is included in Table 1.

3.4 Probability Transition Assignment

The final piece of the puzzle for the MDP model is how to assign transition probabilities to the actions between states. For cyber components, this is almost always set according to domain expertise or security professional advice, as has been a recurring theme throughout this writing. Since this research seeks to determine the importance of accurate expertise on the overall contingency ranking, probabilities will instead be assigned randomly and results will be compared. However, for states in the MDP that reflect having direct control over a power component through the required set of privileges to get access, probability transitions are set to 1. For example, if the attacker gains control of a relay, it is then trivial to open a breaker and cause an outage. There is virtually no mitigation possible between those two events.

3.5 Figures

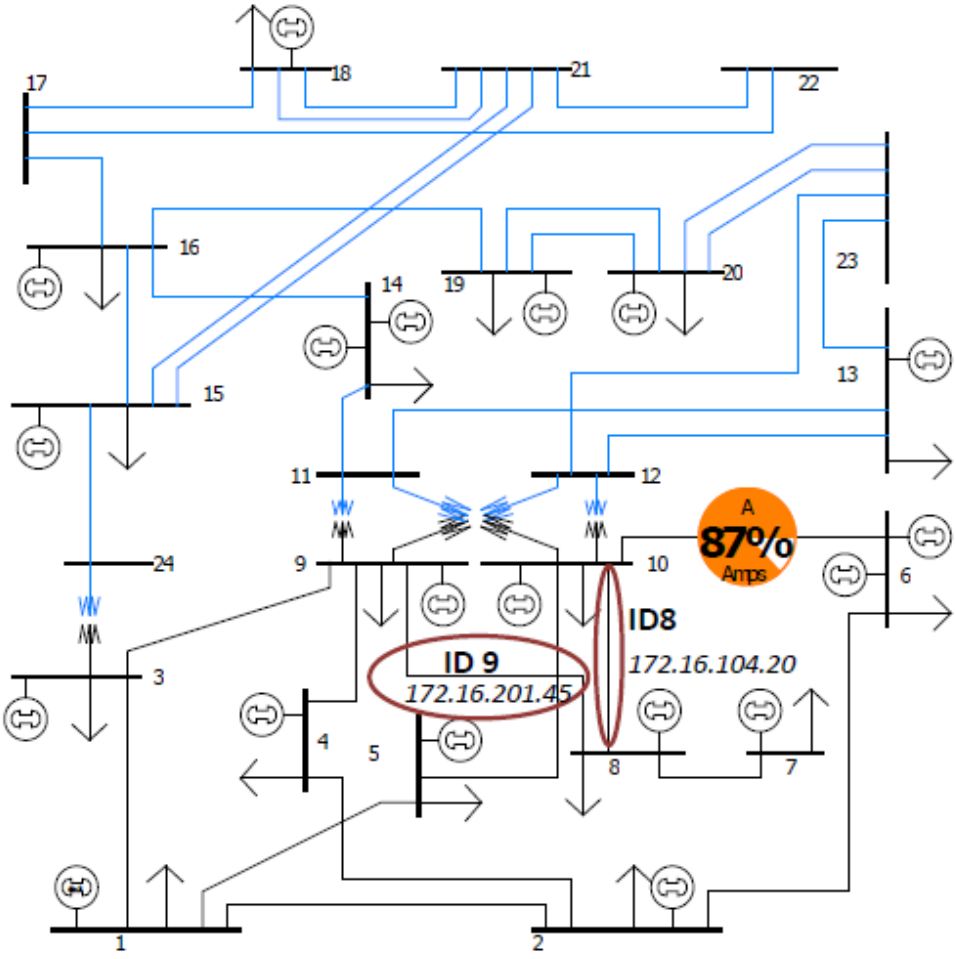

Figure 2: The IEEE 24-bus Reliability Test System

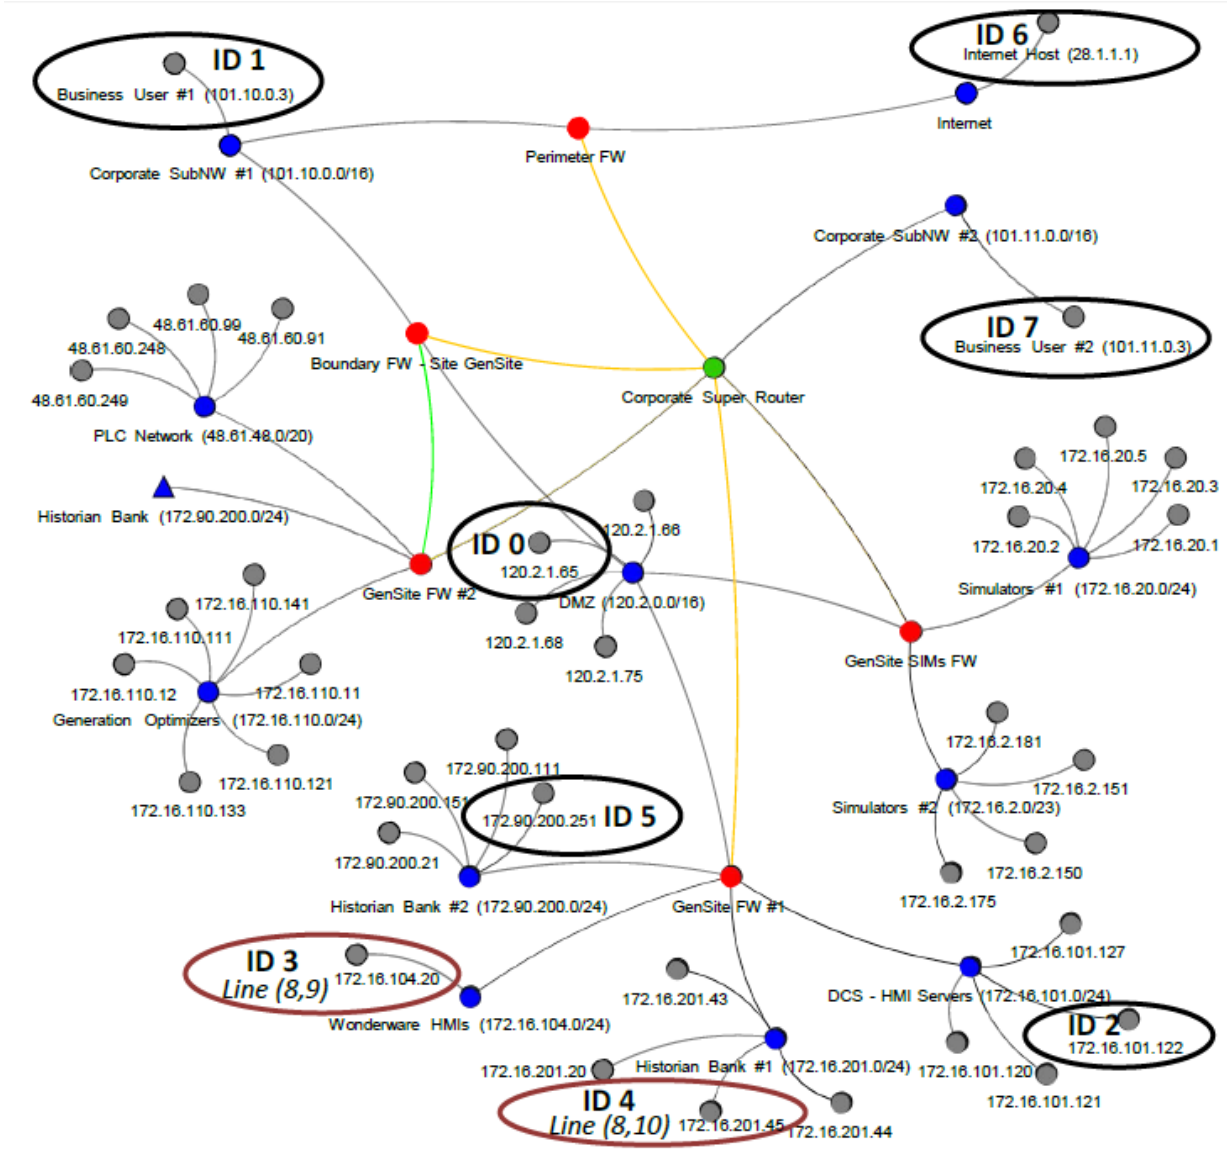

Figure 3: Simulated Cyber Control Network

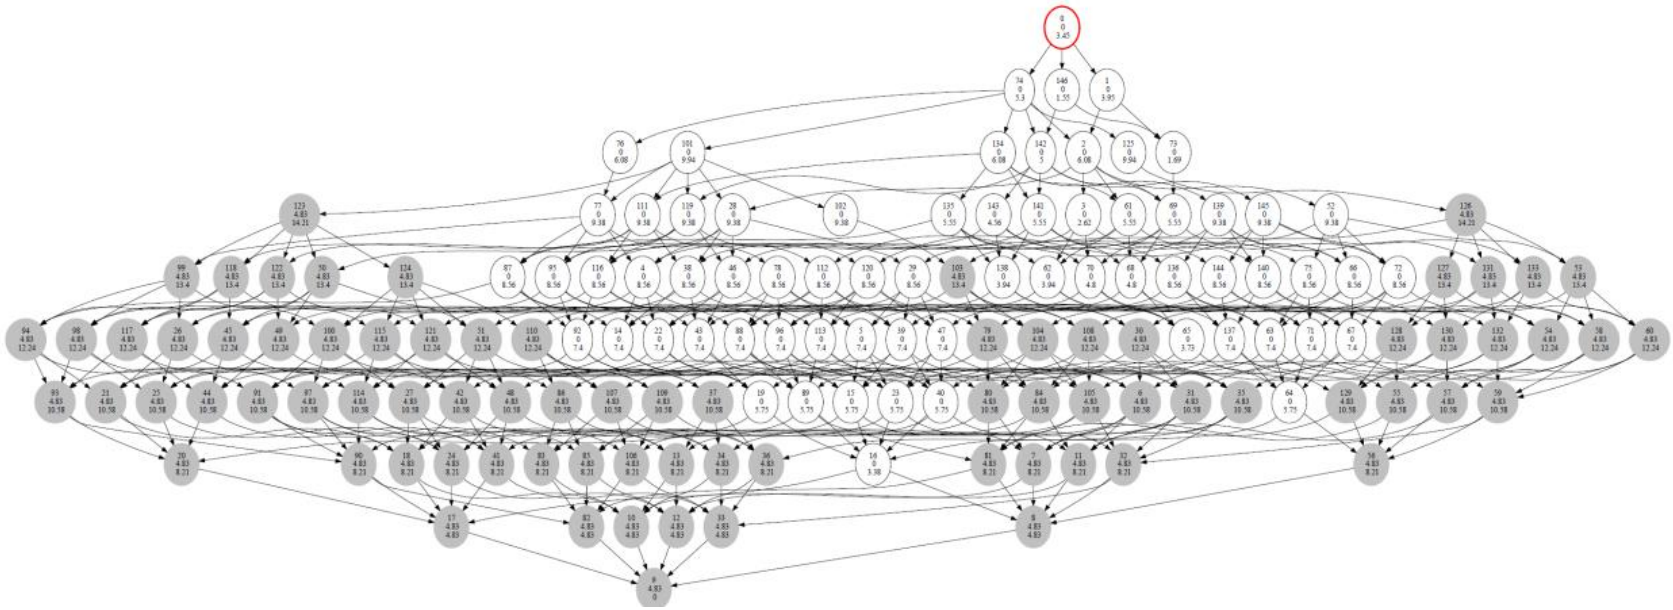

Figure 4: Complete Markov Decision Process

State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets
0	6	1	61	2	610	3	6102	4	61023	5	610234	6	6102349	7	61023495	8	610234957
9	6102349578	10	610234958	11	61023497	12	610234978	13	61023498	14	610235	15	6102354	16	61023547	17	610235478
18	61023548	19	6102357	20	61023578	21	6102358	22	610237	23	6102374	24	61023748	25	6102378	26	610238
27	6102384	28	6103	29	61034	30	610349	31	6103495	32	61034957	33	610349578	34	61034958	35	6103497
36	61034978	37	6103498	38	61035	39	610354	40	6103547	41	61035478	42	6103548	43	610357	44	6103578
45	610358	46	61037	47	610374	48	6103748	49	610378	50	61038	51	610384	52	6104	53	61049
54	610492	55	6104925	56	61049257	57	6104927	58	610495	59	6104957	60	610497	61	6105	62	61052
63	610524	64	6105247	65	610527	66	61054	67	610547	68	61057	69	6107	70	61072	71	610724
72	61074	73	617	74	60	75	61024	76	602	77	6023	78	60234	79	602349	80	6023495
81	60234957	82	602349578	83	60234958	84	6023497	85	60234978	86	6023498	87	60235	88	602354	89	6023547
90	60235478	91	6023548	92	602357	93	6023578	94	602358	95	60237	96	602374	97	6023748	98	602378
99	60238	100	602384	101	603	102	6034	103	60349	104	603495	105	6034957	106	60349578	107	6034958
108	603497	109	6034978	110	603498	111	6035	112	60354	113	603547	114	6035478	115	603548	116	60357
117	603578	118	60358	119	6037	120	60374	121	603748	122	60378	123	6038	124	60384	125	604
126	6049	127	60492	128	604925	129	6049257	130	604927	131	60495	132	604957	133	60497	134	605
135	6052	136	60524	137	605247	138	60527	139	6054	140	60547	141	6057	142	607	143	6072
144	60724	145	6074	146	67												

Table 1: Mapping from Privilege IDs to MDP State IDs

CHAPTER 4: EXPERIMENTATION & RESULTS

The body of work related to cyber-physical contingency analysis for power grids always circles back to the requirement of expert domain knowledge and security knowhow. By nature, it is a tedious and manual process that requires large resources and the benefits are often undermined by the cost. In an MDP model such as this, expertise manifests itself as transition probabilities. Therefore, it is the goal of these experiments to perturb the probabilities assigned to states in the model to determine the importance of and sensitivity to prior knowledge of the attack surface. This will be accomplished by assigning a base value to the transitions in the model as if they were determined by an expert and calculating the security index (value function) accordingly. These probabilities will be chosen randomly out of a set of three possibilities, representing low, medium, and high likelihood of exploitation. These exact values will be 0.2, 0.4, and 0.6 respectively. Once these probabilities have been randomly distributed, the value of each state will be locked in as the base value. Perturbations of the probabilities thereafter should cause changes in the value functions of the MDP states compared to the base values. The change in value of states will be assessed

empirically as well as any change in the contingency ranking results from the base values. In this way, the results will show whether transition probabilities in a MDP modeled after a real power network are highly sensitive to accurate expert analysis or if other factors in the MDP dominate the ranking.

4.1 Interval Value Iteration

At the core of the sensitivity analysis is the Interval Value Iteration (IVI) algorithm [13]. Where traditional MDPs require an exact value for the transition probability function, IVI provides the ability to specify a real interval for each action's likelihood. This is an ideal tool for perturbation because it allows for testing the entire range of probabilities within a given Δp. By testing all probabilities simultaneously, effectively all possible permutations of an MDP, an uncountable amount, are tested in polynomial time. This is incredibly powerful. The authors show that there is a finite set of MDPs that are of interest. Particularly, given an ordering of all the states in the MDP, there is a unique MDP that sends as much probability mass as possible early in the ordering. We will defer to the publication for proofs of order-maximizing MDPs and convergence. As expected, IVI gives a real interval as a result for a state's value function. In fact, it gives optimal policy value

intervals. There are perhaps more sophisticated ways to compare these intervals, but for this study, value intervals were compared by their midpoint to determine ranking. Pseudo-code of IVI can be found in Appendix A.

4.2 Contingency Ranking

The primary way to evaluate whether or not a perturbation of the MDP's probability transition function creates meaningful results is to enumerate the highest ranking contingencies (states with highest value function) and look for changes. In other words, does the attacker's optimal policy change if a different transition function is introduced? Because the attacker is assumed to start outside the network without any privileges, we rank contingencies based on optimal actions from the root. In breadth-first fashion, all actions are enumerated from the root state and the top action is ranked #1. From here, we continue considering the remaining actions from the root state, but we now include all actions reachable from the state in which the system would result if the attacker successfully transitioned from the top contingency. One contingency is ranked per iteration and any new actions reachable given that contingency is taken are added to the list to be ranked. This is a slight break from the algorithm presented in the SOCCA

paper, which ranked all contingencies available before considering the next hop. This continues until all states have been reached and the resulting list is the optimal policy for the attacker. The attacker should take actions in the order listed to maximize the reward, in our case, line outages. Conversely, a system operator should prioritize protecting against actions in the list in order to thwart the attacker most effectively. It should be noted that only cyber contingencies are to be ranked since it is assumed that power contingencies have no option for mitigation and transition with probability 1. To reiterate, the white nodes in the MDP diagram (Figure 4) refer to cyber states while grey nodes denote power states.

4.3 Experiments

Recall that the following experiments were perturbations compared against a base value assigned to states as if an expert has set their probability transition functions. All graphs are located in the Figures subsection of this chapter.

4.3.1 Perturbation of the Base Values

The first experiment conducted was to employ IVI to introduce a Δp for each state's probability function. As the interval is increased outward from the base probability, the possible values associated with the states should significantly change if the probabilities are a dominating factor in the MDP. To test this, the top 10 contingency rankings were compared for each increase in Δp. Additionally, the value intervals returned by IVI in each case were compared and the percent change in value for every state was calculated. The maximum difference of either the interval's lower bound or upper bound from the base value was chosen for the percent change. This data shows how a change in probability directly affects a state's value. An example data set from one iteration of testing is displayed in Table 2.

After carrying out the tests, a surprising result emerged. Increasing the range of probabilities for all states by as much as .3 in either direction caused no change in the contingency rankings at all. Neither the ordering nor the set of contingencies in the top 10 changed at all. Figure 5 shows the average percent difference in value for all states for a given Δp along with the standard deviation. This data set is delimited as 'linear' because the relation

between states with different probability transitions is linear due to the base probabilities being 0.2, 0.4, and 0.6.

4.3.2 Perturbing the Probability Relation

Since no changes were detected in the previous experiment, the next thing to test is the relative relationship between the different probability transition functions. Whereas the probabilities were linearly related before, an exponential relationship was tried. In place of 0.2, 0.4, and 0.6 the new corresponding initial probabilities were set to 0.1, 0.2, and 0.8. Again testing for different values of Δp using IVI, the top 10 rankings were compared to the base rankings and percent change was calculated similarly as before in Figure 6. Please note that base rankings in this context refer to the rankings before any experiments were run, not the beginning of the current experiment.

This time, 9 out of the top 10 rankings were different compared to the base rankings. However, across all Δp using the exponential relationship, there was no change in rankings. This is still a significant change and reflects an entirely different policy that an optimal attacker would pursue. Therefore, the relationship

amongst probability values is an important consideration when constructing the MDP.

4.3.3 Reassigning the Original Probabilities

Another test that was deemed necessary was to detect changes if a whole other assignment of probabilities was given. The base probabilities were assigned randomly out of 0.2, 0.4, and 0.6. For this test, those probabilities were again randomly distributed across all states in a guaranteed unique way from the base assignments. Once more, IVI was utilized to test over increasing values of Δp. The top 10 rankings were again computed and percent change of value functions are given in Figure 7.

Similar to the first experiment, there was no change in the rankings compared to the base rankings across any Δp. Even though the probabilities were scrambled and reordered, there was no consequential change in the MDP's optimal policy. The probabilities within states have a linear relationship as in Experiment 1 so perhaps this is additional evidence for the importance of the relative difference between probability transitions instead of their placement in the MDP.

4.3.4 Perturbing the Discount Factor

The probability transition function is only one aspect that governs the optimal MDP policy. Another variable that may affect change is the discount factor within the value iteration. The discount factor is responsible for diminishing greater rewards that require many actions to attain versus smaller rewards that are more easily within reach. IVI was not used in this test. Instead, the original base probabilities were left constant and the discount factor was varied intermittently between .2 and .9. The base value iteration algorithm used a discount factor of 0.7. Percent change from the base values were again calculated and are included in Figure 8.

Despite seeing the greatest change in the values of states and the widest standard deviation, perturbing the discount factor did not change the top 10 contingency rankings in any way.

4.4 Figures

State		Base Value	V↓(s)	V↑(s)	Δmax	**% Change**
1		2.91	2.21	3.32	0.704	0.242
2		3.39	2.72	3.61	0.67	0.198
3		5.8	5.05	5.82	0.748	0.129
4		2.98	1.62	3.53	1.36	0.456
5		9.38	8.57	8.57	0.811	0.0864
6		8.57	7.41	7.41	1.16	0.135
7		8.57	7.41	7.41	1.16	0.135
8		7.41	5.75	5.75	1.66	0.224
9		5.75	3.38	3.38	2.37	0.412
10		7.41	5.75	5.75	1.66	0.224
11		8.57	7.41	7.41	1.16	0.135
12		7.41	5.75	5.75	1.66	0.224
13		9.95	9.38	9.38	0.569	0.0571
14		9.38	8.57	8.57	0.811	0.0864
15		9.38	8.57	8.57	0.811	0.0864
16		8.57	7.41	7.41	1.16	0.135
17		7.41	5.75	5.75	1.66	0.224
18		8.57	7.41	7.41	1.16	0.135
19		9.38	8.57	8.57	0.811	0.0864
20		8.57	7.41	7.41	1.16	0.135
21		9.95	9.38	9.38	0.569	0.0571
22		5.47	4.61	5.32	0.856	0.156
23		4.14	3.05	3.99	1.09	0.263
24		8.57	7.41	7.41	1.16	0.135
25		7.41	5.75	5.75	1.66	0.224
26		4.32	3.1	3.57	1.22	0.283
27		9.38	8.57	8.57	0.811	0.0864
28		8.57	7.41	7.41	1.16	0.135
29		5	3.99	4.6	1.01	0.202
30		5.47	4.61	5.32	0.856	0.156
…		…	…	…	…	…

Table 2: Example Data Set from Experiment 1
Δp = .1

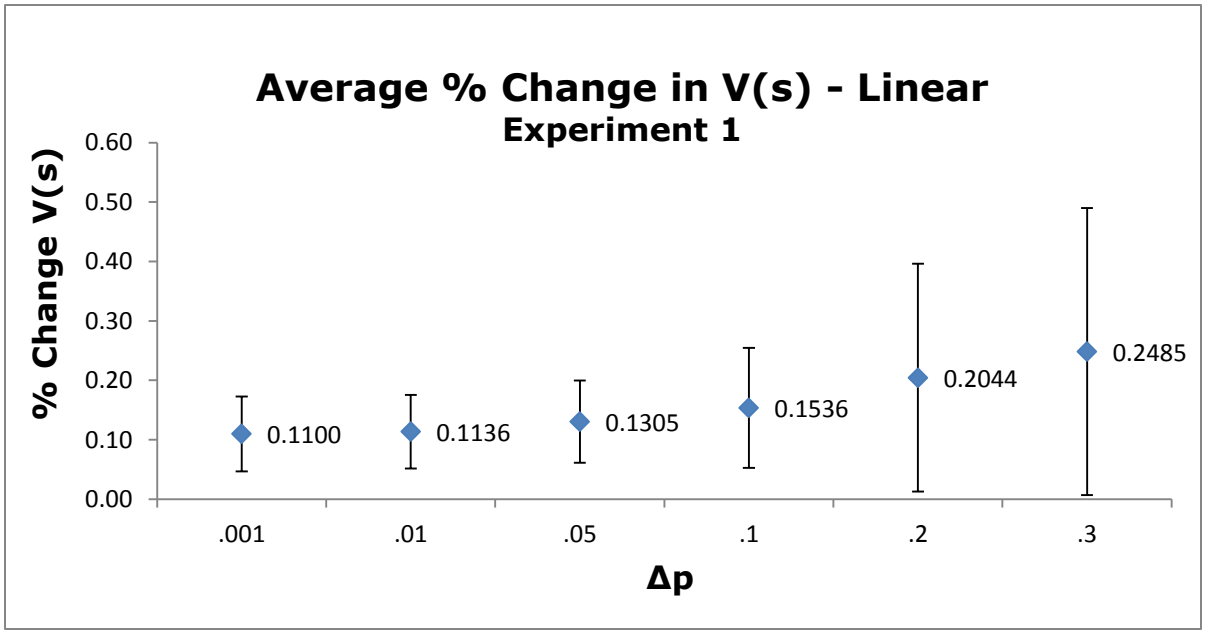

Figure 5

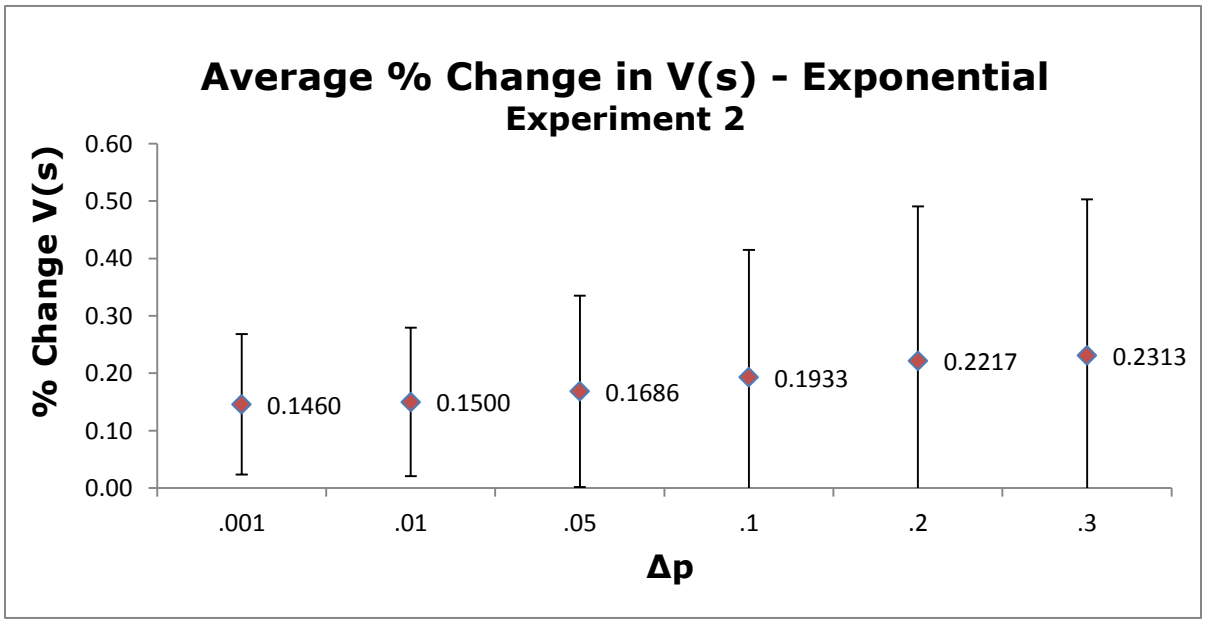

Figure 6

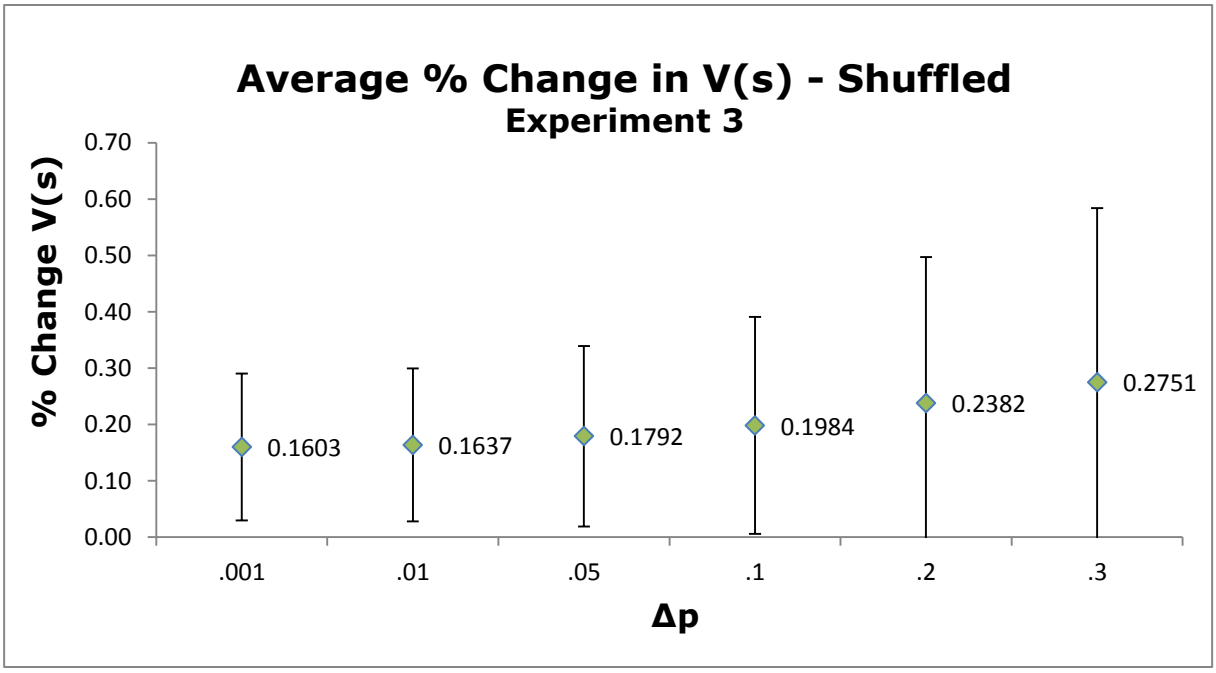

Figure 7

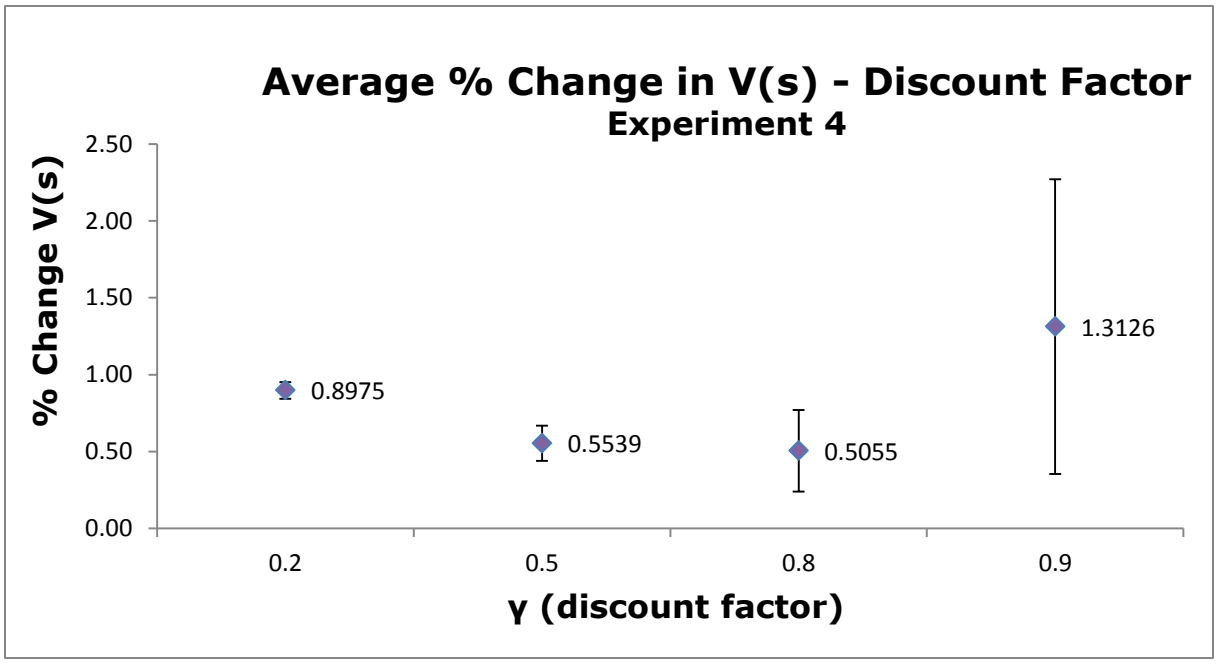

Figure 8

Top 10 Base Contingency Rankings		
Start State	End State	V(s)
0	74	4.99
74	101	10.34
74	125	10.34
101	28	9.94
101	77	9.94
101	102	9.94
101	111	9.94
101	119	9.94
28	4	9.38
77	4	9.94

Table 3

Top 10 Contingency Ranking Changes vs. Base	
Across all Linear	0
From Base to Exponential	9
Across All Exponential	0
From Base to Shuffled Linear	0
Across All Shuffled	0
Across All Change in Discount Factor	0

Table 4

CHAPTER 5: CONCLUSIONS

Towards the objective of evaluating the sensitivity of a prototypical power network MDP to the effects of probability transitions, the results show conclusive evidence that the role of expert knowledge is not as integral to contingency ranking as previously thought. Despite augmenting the security index of MDP states by as much as 130% on average with standard deviation approaching the average itself, contingency rankings were shown to not be affected at all. The discount factor of value iteration was also determined to be inconsequential albeit being the most influential in terms of value change. The only noticed change in contingency rankings occurred when greater disparity was placed in the relative values of available probability transitions, though the placement of the transitions themselves seems insignificant.

This seems to imply that the native structure of a power network is governed by the only remaining variable, the assignment of rewards amongst states. This is reasonable when considering that power control networks are typically structured such that there are only a few hops between the outside Internet and components that directly affect line flow. This is also in keeping with the radial

topology of the control network. If the initial few hops are breached, the inner workings of the network are open to exploration and exploit. Therefore, it is imperative to maintain a strong outer shell and place defense mechanisms early in the attack chain. Deriving the most vulnerable assets is an ongoing problem, but it seems that focus ought to be shifted towards accurate reward functions and broad probability transition schemes in place of individual probability scrutiny in the scope of MDP modeling.

5.1 Future Work

Going forward with this reasoning, it seems prudent to perform experiments using varying reward functions to identify the most accurate measure of contingency ranking. Additionally, the incorporation of negative rewards for attack failures may provide additional insight into the behavior of the system. Furthermore, probability transition functions that model tight exterior mitigation may be equally interesting.

As for extension of the SOCCA framework into the future, alert correlation may provide an avenue to enable real-time threat monitoring and contingency warnings. Works such as [15] present

novel approaches to alert correlation through attack trees. By mapping intrusion detection or other network management alerts to assets in an attack chain as we have demonstrated using MDPs, the current state of the network can be approximated in good confidence and re-evaluation of contingencies can take place quickly and responsively. A working prototype of this type of analysis has been discussed and implemented in a related work to SOCCA called SCPSE [16].

REFERENCES

[1] A. Monticelli, *State Estimation in Electric Power Systems: A Generalized Approach*. Kluwer Academic Publishers, 1999.

[2] R.S. Maciel, A. Padilha-Feltrin, E. Righeto, "Substitution-Newton-Raphson Method for the Solution of Electric Network Equations," *Transmission & Distribution Conference and Exposition: Latin America, 2006. TDC '06. IEEE/PES*

[3] J. Deuse, K. Karoui, A. Bihain, J. Dubois, "Comprehensive approach of power system contingency analysis," *Power Tech Conference Proceedings, 2003 IEEE Bologna*

[4] S. Zonouz, C.M. Davis, K.R. Davis, R. Berthier, R.B. Bobba, W.H. Sanders, "SOCCA: A Security-Oriented Cyber-Physical Contingency Analysis in Power Infrastructures," *Smart Grid, IEEE Transactions on* , vol.5, no.1, pp.3,13, Jan. 2014

[5] D. M. Nicol, W. H. Sanders, S. Singh, and M. Seri, "Usable global network access policy for process control systems," *IEEE Security and Privacy*, vol. 6, pp. 30–36, 2008

[6] J. Glover, M. Sarma, T. Overbye, and T. Overbye, *Power system analysis and design*. Thomson, 2008

[7] R. Bellman, *A Markovian Decision Process*. Journal of Mathematics and Mechanics 6, 1957

[8] T. M. Chen, S. Member, J. C. Sanchez-aarnoutse, and J. Buford, "Petri net modeling of cyber-physical attacks on smart grid," *Smart Grid IEEE Transactions on*, vol. 2, no. 99, pp. 1–9, 2011.

[9] F. Pasqualetti, F. Dörfler, and F. Bullo, "Cyber-physical attacks in power networks: Models, fundamental limitations and monitor design," *CoRR*, vol. abs/1103.2795, 2011

[10] X. Cao, H. Chen, "Perturbation realization, potentials, and sensitivity analysis of Markov processes," *Automatic Control, IEEE Transactions on* , vol.42, no.10, pp.1382,1393, Oct 1997

[11] S. Ross, M. Izadi, M. Mercer, and D. Buckeridge, "Sensitivity analysis of POMDP value functions", *Proc. IEEE Int. Conf. Mach. Learning Appl.*, pp.317 -323 2009

[12] S. Kalyanasundaram, E. K. P. Chong, and N. B. Shroff, "Markov decision processes with uncertain transition rates: sensitivity and robust control," *Decision and Control, 2002, Proceedings of the 41st IEEE Conference on* , vol.4, no., pp.3799,3804 vol.4, 10-13 Dec. 2002

[13] R. Givan, S. Leach, and T. Dean.. *Bounded Parameter Markov Decision Processes*. Technical Report. Brown University, Providence, RI, USA. 1997.

[14] R. T. S. T. F. of the Application of Probability Methods Subcommittee, "IEEE reliability test system," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-98, no. 6, pp. 2047–2054, Nov. 1979

[15] S. Roschke, F. Cheng, and C. Meine, "A new alert correlation algorithm based on attack graph," *Proceedings of the 4th international conference on Computational intelligence in security for information systems* (CISIS'11)

[16] S. Zonouz, K.M. Rogers, R. Berthier, R.B. Bobba, W.H. Sanders, T.J. Overbye, "SCPSE: Security-Oriented Cyber-Physical State Estimation for Power Grid Critical Infrastructures," *Smart Grid, IEEE Transactions on* , vol.3, no.4, pp.1790,1799, Dec. 2012

APPENDIX A: INTERVAL VALUE ITERATION ALGORITHM

$IVI_{\updownarrow\,\mathrm{opt}}(V_{\updownarrow})$

we assume that $V_{\updownarrow}$ is represented as:
\\ $V_{\downarrow}$ *is a vector of n real numbers giving lower-bounds for states* q_1 *to* q_n
\\ $V_{\uparrow}$ *is a vector of n real numbers giving upper-bounds for states* q_1 *to* q_n

{ Create O, a vector of n states for holding a permutation of the states q_1 to q_n

first, compute new lower bounds
O = sort_increasing_order(q_1,...,q_n,$<_{\mathrm{lb}}$); \\ $<_{\mathrm{lb}}$ *compares state lower-bounds*
VI-Update($V_{\downarrow}$, O);

second, compute new upper bounds
O = sort_decreasing_order(q_1,...,q_n,$<_{\mathrm{ub}}$); \\ $<_{\mathrm{ub}}$ *compares state upper-bnds*
VI-Update($V_{\uparrow}$, O)}

==========

\\ VI-Update(v, o) *updates v using the order-maximizing MDP for o*
\\ *o is a state ordering—a vector of states (a permutation of* q_1,...,q_n*)*
\\ *v is a value function—a vector of real numbers of length n*

VI-Update(v, o)
{ Create F_a, a matrix of n by n real numbers for each action a

\\ *the next loop sets each* F_a *to describe a in the order-maximizing MDP for o*
for each state p and action a {
used = $\sum_{\text{state } q} F_{\downarrow p,\,q}(a)$;
remaining = 1 – used;
\\ *distribute remaining probability mass to states earlier in ordering*
for i=1 to n { \\ i *is used to index into ordering o*
min = $F_{\downarrow p,\,o(i)}(a)$;
desired = $F_{\uparrow p,\,o(i)}(a)$;
if (desired <= remaining)
then $F_a(p,o(\mathrm{i}))$ = min+desired;
else $F_a(p,o(\mathrm{i}))$ = min+remaining;
remaining = max(0,remaining-desired)}}

\\ F_a *now describes a in the order-maximizing MDP w/respect to O,*
\\ *finally, update v using a value iteration-like update based on F'*
for each state p
$v(p) = \max_{a \in A} \left[R(p) + \gamma \sum_{\text{state } q} F_a(p,q)\, v(q) \right]$

Figure 9: Pseudo-code of the IVI algorithm [13]