

**PURDUE UNIVERSITY  
GRADUATE SCHOOL  
Thesis/Dissertation Acceptance**

This is to certify that the thesis/dissertation prepared

By Paul Wood

Entitled

IMPROVING THE RESILIENCE OF CYBER-PHYSICAL SYSTEMS UNDER STRATEGIC ADVERSARIES

For the degree of Doctor of Philosophy

Is approved by the final examining committee:

Saurabh Bagchi

Chair

Milind Kulkarni

Charles A. Bouman

Alefiya Hussain

To the best of my knowledge and as understood by the student in the Thesis/Dissertation Agreement, Publication Delay, and Certification Disclaimer (Graduate School Form 32), this thesis/dissertation adheres to the provisions of Purdue University's "Policy of Integrity in Research" and the use of copyright material.

Approved by Major Professor(s): Saurabh Bagchi

Approved by: Venkataramanan Balakrishnan

Head of the Departmental Graduate Program

Date

IMPROVING THE RESILIENCE OF  
CYBER-PHYSICAL SYSTEMS UNDER STRATEGIC ADVERSARIES

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Paul Wood

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

August 2016

Purdue University

West Lafayette, Indiana

To my family who provided me with the tools to start this journey.  
To my girlfriend Vy who provided me with the endurance to complete it.

## ACKNOWLEDGMENTS

I would like to acknowledge the many discussions and feedback that I have received throughout my time at Purdue. First, I would like to thank my advisor Dr. Saurabh Bagchi who provided me with pivotal feedback throughout my thesis process. His thought-provoking criticisms and suggestions immensely strengthened my technical writing skills, idea presentation, and the technical quality of the solutions in this thesis. Second, I would like to thank Dr. Alefiya Hussain who supported me as a visiting research assistant at USC/ISI. Her mentorship and guidance helped tailor my thesis toward concrete cyber-physical system problems. I would also like to thank Tom Nudell and Dylan Shiltz for providing technical feedback on the power and control system aspects of CPS and Chris Gutierrez for many discussions about my thesis and goals at Purdue. I would also like to thank the other members of the Dependable Computing Systems Lab (DCSL) for broad feedback and support in all aspects of my thesis.

This thesis would not be possible without the financial support that I received from several sources throughout my time at Purdue: the Missile Defense Agency supported my graduate studies for many semesters in System of Systems Analytics, the Andrews Fellowship allowed me to pursue my unique thesis, the DETER project supported my time at USC/ISI, and the Frederic Miller scholarship and REACH fellowship supported many of the additional costs of pursuing my degree.

I am privileged to have worked with such a great group of individuals.

## TABLE OF CONTENTS

	Page
LIST OF TABLES . . . . .	x
LIST OF FIGURES . . . . .	xi
ABSTRACT . . . . .	xviii
1 INTRODUCTION . . . . .	1
1.1 Motivation . . . . .	1
1.2 Contributions and Outline . . . . .	3
1.2.1 Energy CPS Models and Attack Strategies . . . . .	3
1.2.2 Energy CPS Operation . . . . .	4
1.2.3 CPS Attacks and Market Manipulations . . . . .	5
1.2.4 Low Cost DDoS Defense . . . . .	5
1.3 Work Publication . . . . .	6
2 ENERGY-BASED CPS MODELS AND HIGH-LEVEL ATTACK STRATEGIES . . . . .	8
2.1 A Profit-Oriented Model for Energy CPS . . . . .	8
2.1.1 Introduction . . . . .	8
2.1.2 Background . . . . .	10
2.1.3 Related Work . . . . .	14
2.1.4 Overview . . . . .	16
2.1.5 Approach . . . . .	17
2.1.6 Experimental Model . . . . .	25
2.1.7 Results . . . . .	30
2.1.8 Security Strategy Analysis . . . . .	34
2.1.9 Conclusion . . . . .	35
2.2 High-Level Adversarial Strategies . . . . .	37

	Page
2.2.1 Introduction . . . . .	37
2.2.2 Model Overview . . . . .	39
2.2.3 Attacker Strategy . . . . .	40
2.2.4 Defense . . . . .	42
2.2.5 Experimentation . . . . .	45
2.2.6 Conclusion . . . . .	52
2.3 Model and Attacks Conclusions, Insights, and Future Work . . . . .	53
3 ENERGY CPS OPERATION . . . . .	55
3.1 A Framework for Evaluating the Resilience of Dynamic Real-Time Market Mechanisms . . . . .	55
3.1.1 Introduction . . . . .	55
3.1.2 Dynamic Real-Time Market Mechanisms . . . . .	60
3.1.3 Modeling Latency in Communication Networks . . . . .	64
3.1.4 Resilience Metric . . . . .	68
3.1.5 Experimental Setup . . . . .	69
3.1.6 Experimental Results . . . . .	73
3.1.7 Conclusions . . . . .	77
3.2 Dynamic Pricing for Smart Grids in the Presence of Non-Linear Net- work Conditions . . . . .	78
3.2.1 Introduction . . . . .	78
3.2.2 Background . . . . .	82
3.2.3 Real-Time Communication . . . . .	83
3.2.4 Related Work . . . . .	85
3.2.5 Existing DSM Techniques . . . . .	85
3.2.6 Existing Latency Studies . . . . .	86
3.2.7 Solution . . . . .	87
3.3 Networking and DSM . . . . .	91
3.3.1 Latency Influences in DSM . . . . .	91
3.4 Experimental Results . . . . .	94

	Page
3.4.1 Experimental Setup . . . . .	94
3.4.2 Load and Generation Model . . . . .	95
3.4.3 Experiment 1: Analysis under Normal Conditions . . . . .	96
3.4.4 Experiment 2: Unpredictable Step Change . . . . .	98
3.4.5 Experiment 3: Sensitivity to $\beta$ . . . . .	99
3.4.6 Experiment 4: Networking Impacts . . . . .	101
3.4.7 Discussion . . . . .	102
3.4.8 Conclusion . . . . .	104
3.5 Technical Market Conclusion . . . . .	104
VOLUME 2	
4 MARKET MANIPULATION . . . . .	105
4.1 Introduction . . . . .	105
4.2 Preliminaries . . . . .	108
4.2.1 Electric Power Grids . . . . .	108
4.2.2 Power Markets . . . . .	109
4.2.3 Profit Manipulation . . . . .	111
4.3 Attack and Defense Strategy . . . . .	111
4.3.1 Players and Definitions . . . . .	113
4.3.2 Attacker's Incentive . . . . .	113
4.3.3 Defensive Maneuvers . . . . .	114
4.3.4 Defensive System Overview . . . . .	115
4.4 Multiple Knowledge Levels . . . . .	117
4.4.1 Multiple Underlying Games . . . . .	117
4.4.2 Perfect Information Game . . . . .	118
4.4.3 Imperfect Attack Strategies . . . . .	118
4.4.4 Multiple Defender Optimization . . . . .	119
4.5 Experimentation . . . . .	120
4.5.1 Experimental Setup . . . . .	120

	Page
4.5.2 Experiment 1: The Attacker's Incentive . . . . .	121
4.5.3 Experiment 2: Collaborating Defenders . . . . .	122
4.6 Related Work . . . . .	124
4.7 Conclusion . . . . .	126
5 NETWORK ATTACKS IN MARKETS . . . . .	128
5.1 Introduction . . . . .	128
5.2 Background and Market Operations . . . . .	132
5.2.1 Real-time Markets . . . . .	132
5.2.2 Demand-Side Management . . . . .	132
5.2.3 Real-Time Communication . . . . .	133
5.3 Strategic Adversary . . . . .	135
5.3.1 Strategy Summary . . . . .	135
5.3.2 Capabilities and Resources . . . . .	135
5.3.3 Strategy Definition . . . . .	137
5.3.4 Price Manipulation . . . . .	137
5.3.5 Attack Strategy . . . . .	139
5.3.6 Integrity Attacks . . . . .	140
5.4 Defender Strategies . . . . .	141
5.4.1 Moving Target Defense . . . . .	142
5.4.2 Detection via Deception . . . . .	143
5.5 Experimental Setup . . . . .	143
5.5.1 Overview of RTP Controller . . . . .	144
5.5.2 Load and Generation Model . . . . .	144
5.5.3 Real-World Dataset . . . . .	146
5.6 Experimental Results . . . . .	146
5.6.1 Experiment I: Baseline Profits . . . . .	147
5.6.2 Experiment II: Impact of DoS Attacks . . . . .	147
5.6.3 Experiment III: Impact of Integrity Attacks . . . . .	149



	Page
5.6.4 Experiment IV: Defensive Strategies . . . . .	150
5.6.5 Return on Investment . . . . .	152
5.7 Discussion . . . . .	153
5.7.1 Grid Dynamics in Real-Time . . . . .	153
5.7.2 Attribution for Attacks . . . . .	154
5.8 Conclusion . . . . .	154
6 DEFENSES . . . . .	156
6.1 Introduction . . . . .	156
6.2 Background and Related Work . . . . .	159
6.2.1 Traditional Defenses . . . . .	159
6.2.2 Overlays and Moving Target Defenses . . . . .	159
6.3 DoSE Overview . . . . .	160
6.3.1 Threat Model . . . . .	160
6.3.2 Workflow of DoSE . . . . .	160
6.3.3 Client-Relay Assignment . . . . .	161
6.3.4 Defense Against Some Obvious Adversaries . . . . .	162
6.3.5 Cost Minimization . . . . .	162
6.4 Detailed Design of DoSE . . . . .	163
6.4.1 Relays . . . . .	163
6.4.2 Client-Relay Assignment Infrastructure . . . . .	163
6.4.3 Client-Relay Assignment Strategy Overview . . . . .	167
6.4.4 Formal Relay Assignment Strategy . . . . .	168
6.4.5 Optimizing Cost . . . . .	169
6.5 Experimental Results . . . . .	169
6.5.1 Management Service Overhead Analysis . . . . .	170
6.5.2 Agent-Based Simulator . . . . .	170
6.5.3 Experiment 1: Single Adversary . . . . .	171
6.5.4 Experiment 2: Streaming Attack . . . . .	172

	Page
6.5.5 Experiment 3: Attackers Present at Startup . . . . .	173
6.6 Result Analysis . . . . .	175
6.6.1 Sensitivity to CRPA . . . . .	175
6.6.2 Quantitative Comparison with Epiphany and Speak-up . . .	176
6.6.3 Comparison with MOTAG . . . . .	178
6.7 Costs and Discussion . . . . .	180
6.8 Conclusion . . . . .	183
7 CONCLUSION AND FUTURE WORK . . . . .	184
7.1 Conclusion . . . . .	184
7.2 Future Work . . . . .	185
7.2.1 Multi-Round Attacker/Defender Games . . . . .	185
7.2.2 Stateful Real-Time Attacks . . . . .	186
7.2.3 Client-Side DoSE . . . . .	186
7.2.4 Algorithm Resilience . . . . .	186
A SIMULATION ENVIRONMENT . . . . .	188
A.1 Game Theoretic Model for CPS . . . . .	188
A.2 Agent-Based Model . . . . .	188
A.2.1 Technical Markets . . . . .	189
A.2.2 Technical Markets and Game Theory . . . . .	189
A.2.3 Real-Time Market Attacks . . . . .	189
A.2.4 DoSE Model . . . . .	189
LIST OF REFERENCES . . . . .	190
VITA . . . . .	199

## LIST OF TABLES

Table	Page
3.1 Example parameters used in Figure 3.9 . . . . .	95
4.1 Example Impact Matrix (IM) . . . . .	115
5.1 List of Symbols . . . . .	135
5.2 Parameter distributions used in Experiments, respectively for the Consumer (C) and the Generator (G) . . . . .	145

## LIST OF FIGURES

Figure	Page
2.1 High level layout of the natural gas infrastructure. . . . .	11
2.2 High level layout of the electric infrastructure. . . . .	12
2.3 Flow graph of the combined natural gas and electric infrastructure. . .	21
2.4 The graph shows an example scenario where each asset is owned by a different actor A-F. The three suppliers A, B, D have unit cost of production \$3, \$3, and \$9 respectively and the end customer has a fixed unit price of \$11. At each edge, there is a fixed unit cost and a marginal unit cost. The fixed cost is a parameter of the model, while the marginal cost is calculated by observing the system's response to constricting flow through that particular edge. In this case, actors A and B are in direct competition, so no unilateral price movement is possible. C, however, is in a position to mark up the price an additional \$6 because it is in competition with D. E and F are in series and may mark up to the remaining purchase price at the customer, cumulatively. A fair split of \$1 for E and \$1 for F is shown as an example. . . . .	24
2.5 A flow model is created for six Western US states for both an electric (a) and natural gas (b) infrastructure. . . . .	26
2.6 This chart shows the marginal change in the impact for incremental reductions in capacity for four edges in the system model. The incremental steps represent the costs of the marginal generation that must be utilized to deliver energy to the customer. . . . .	31
2.7 This chart shows the impact of objective independence in ten edges due to a failure in the model. <b>Combined</b> indicates a single actor model where as <b>Independent</b> indicates two actors acting independently. . . . .	33
2.8 This chart compares the normalized financial impact (left) to the shortage (right) caused during a complete failure of an edge in the system. The impact analysis could be different based on which objective is chosen, profit maximization or shortage minimization. . . . .	34

Figure	Page
2.9 Impact of edge failure compared with its transport capacity. A lack of a clear increasing trend suggests that this approach is non-optimal in attacking. (b) shows how the impact of the two points progresses differently as capacity is reduced. Since the impact is low for a large capacity loss in the circle point, based on the slope of the line, the attack has less overall impact on the system. . . . .	36
2.10 Impact of edge failure compared with its optimized flow. The high flow edges, when attacked, result in the most redirection of energy in the system. (b) shows how the alternative energy cost may be much higher per watt for some sources than others, the greedy approach will not work as show in (a). . . . .	37
2.11 The total gain and loss in the system, as the sum across impacts felt by all actors, increase as the number of actors in the system increase up to a point of saturation. The sum of the gain and negative loss remain constant. . . . .	46
2.12 This figure shows the profitability of the strategic adversary versus the amount of knowledge (inverse of noise) that it has about the system. As the noise increases, the profitability decreases. Additionally, as the number of actors increases, the profitability of the SA also increases because of profit opportunities. . . . .	48
2.13 This compares the profit of attack for a 6-actor system. The SA anticipated returns, based on the noisy model, do not decay with knowledge level. This means that if the SA is overconfident, the observed returns will be much less than anticipated. . . . .	49
2.14 The effectiveness of a defense is graded by its impact reduction in ground truth versus the knowledge level of the defender, modeled as noise added to the ground truth. As the number of actors increases, the effectiveness of the defense decreases due to misaligned incentives and a lack of pooled defensive budgets. . . . .	50
2.15 The impact of collaboration is measured by allowing the actors to share in defensive costs. When the costs are shared, more effective investments can be made. . . . .	51
2.16 Collaboration allows for actors to improve their defenses. In this case, the system-wide defensive investment is fixed as the number of actors increases, resulting in reduced benefit of collaboration as the number of actors increases and their individual budgets dwindle. . . . .	52

Figure	Page
3.1 Each agent in the physical network—generators and flexible loads—is associated with a client in the communication network, these clients send and receive information to and from the server at the independent system operator (ISO) who operates the market. . . . .	60
3.2 Relevant time scales for distributed market mechanisms . . . . .	65
3.3 The system responds to a worst-case single-generator failure. The market minimizes residual power in two identical experiments, but one has a communication link outage to a single client—the one that causes the most impact. . . . .	73
3.4 The mean latency for each client is increased. The higher latency values reduce the rate of communication, and the convergence speed of the system for different latencies is shown. . . . .	75
3.5 The variance in latency is increased for a small mean value. After the system apparently stabilizes, the response begins to oscillate. . . . .	76
3.6 Communication pattern: Negotiation phase and actuation phase. . . .	84
3.7 Faults: When the communication channels are interrupted, the algorithm operation becomes blocked (left). With a recovery mechanism, a new price and power level can be used to continue operations (right). . . . .	91
3.8 Asynchronous state traversal from $x(t_0)$ to $x(t_1)$ via several potential paths, each a function of communication latency. . . . .	93
3.9 The supply (generator) and demand (consumer) are shown for varying market prices. The residual power shown is the imbalance between the generator's output and the consumer's input. Since residual power can disrupt the grid and lead to inefficient generation, the goal is to optimize the market price such that the residual power is minimized. . . . .	96
3.10 The $\lambda$ price values are shown for a particular one-hour scenario. The optimal price is constantly changing, and the online method is attempting to match it. The time-of-use (TOU) method provides pricing windows with perfect prediction. The online method is much closer to the optimal price than the TOU approach due to increased granularity of control. . .	97
3.11 The residual power in the market must be absorbed by automatic generation control. Less residual power indicates optimality in the pricing method. The online method is able to outperform the TOU method when changes occur in the system due to increased time-precision, reducing residual power by 64%. When the system is not changing rapidly, at the 30-minute mark, both approaches are comparable. . . . .	98

Figure	Page
3.12 A step load is added and removed at 20 and 40 minutes respectively. The TOU method, calculated a priori, is unable to adapt to the change in the system. The online method lags the optimal method during the step transient, but it is able to adapt to the changing market conditions. . . . .	99
3.13 The residual power is shown for a step change in load. The online method responds to the transient to curb residual power and outperforms the a prior TOU method by 85%. . . . .	100
3.14 Different $\beta$ values for Algorithm 1 are plotted for a step change at time 200-400 seconds. The high values of $\beta$ improve adaptability to transients while the low values resist drastic changes to the price. . . . .	101
3.15 The average residual power in the market decreases with higher $\beta$ values during the transient period. Without transients, the $\beta$ parameter has little impact on performance. . . . .	101
3.16 This figure shows the performance of the system with a varying number of clients experiencing latency conditions. As more clients are exposed to latency, the system is unable to adapt well to transients. Small populations are easily accounted for by the online adaptability of the algorithm, but the larger populations place a high reliance on the load-estimation function. . . . .	103
4.1 The residual power (RP) and market price ( $\lambda$ ) are shown for an example grid scenario based on prior market solution work, later described in Section 4.5.1. The market experiences a demand surge in $\bar{P}_i$ for consumers at $t=100$ s followed by a reduction at $t=250$ s. During the surge, residual power spikes until the market price is corrected. . . . .	110
4.2 The profitability for a generator is shown for two attack scenarios. At $t=50$ s a DoS attack is launched on the communication link connecting the generator to the market (self-attacked) or another market player (others attacked), and it lasts until $t=200$ s. . . . .	112
4.3 In the overall system flow, set of defenders and a strategic adversary each have a view of the system and its market. The system is exposed to a physical scenario and analyzed for a set of potential targets via a market mechanism simulation. From this simulation, the profitability of each market player is captured in a set of impact matrices (IM). Each market player has an independently calculated IM and thus a different defense strategy that can be rectified via collaboration. . . . .	116
4.4 The simulated network has four top-tier links, twelve mid-tier links, and one hundred twenty leaf links. . . . .	121

Figure	Page	
4.5	The strategic adversary attempts to maximize her incentive by disrupting network links. In the graph on the left, the adversary is disrupting leaf-links in the communication topology. In the middle graph, the adversary is disrupting mid-tier links, and in the graph on the right, the top tier links are disrupted. In each case, the attacker's incentive is maximized for the given targets attacked. The strategy shown is maximized from (4.5) and compared to the mean of a random target selection. . . . .	122
4.6	The attacker's incentive is reduced by defensive investments in communication links. As the number of links attacked increases, the number of links defended also increases. The effectiveness of the defense, however, is reduced by imperfect knowledge levels among the defenders ( $\sigma = 0.1$ ). Each line represents a different amount of aggregate defense budget, relative to the number of links attacked. . . . .	123
4.7	This figure shows the cumulative reduction in attacker's incentive for different defender knowledge levels when 75 targets are attacked and defended. Decreased knowledge levels (high $\sigma$ values) results in ineffective defense. As defenders are unwilling and unable to collaborate on defensive investments, the system suffers overall from poor defensive strategies. . . . .	124
5.1	The supply (generator) and demand (consumer) are shown for varying market prices. The residual power shown is the imbalance between the generator's output and the consumer's input. Since residual power can disrupt the grid and lead to inefficiencies in the energy use, the goal is to optimize the market price such that the residual power is minimized. . . . .	145
5.2	The market price during optimal baseline operation is shown. The charge and discharge markers indicate the adversary's optimal charging strategy. . . . .	148
5.3	The power of attack and the gain for 20 targets is shown for the day. The $D_j$ term becomes saturated at high market prices due to output saturation at the largest market players. . . . .	149
5.4	The market price is shown when the attacker implements a DoS attack strategy on 20 targets. . . . .	149
5.5	The market price is shown when the attacker implements a integrity attack strategy on 20 targets. During the charging phase, consumers are misled into conserving power, and during discharge, consumers are misled into over-purchasing power, and this results in price increases and decreases for the adversary to leverage. . . . .	151
5.6	The reduction of the adversary's attack-induced profit is shown as her information about the targets decreases. Errors in target value effectively reduce the profit of the adversary. . . . .	152



Figure	Page
5.7 The reduction of the adversary's attack-induced profit is shown as targets are swapped by the defender. Since the swaps are optimized on target impact, there is a diminishing return on investment for swapping all of the assets. The first swap protects a large generator with the most impact.	153
6.1 Overview of Infrastructure deployed as part of DOSE: The attackers and clients exist on the Internet and are attempting to access a protected service that is either located in a data center or on a small business network (SBN).	161
6.2 The client interacts with static files stored on the CDN to retrieve an identity for obtaining future relay assignments in the event of an attack. The client selects a random puzzle from a set on the CDN and then does a proof-of-work to solve the puzzle and retrieve an initial Client ID and relay assignment to begin accessing the protected service.	166
6.3 Percentage of Failed Transactions with a single attacker: The attack begins at 80 seconds and is progressively mitigated with each relay-creation cycle until the attacker is identified and neutralized at 3.9 minutes. The relay power-on time is 40 seconds, the approximate width of the steps.	172
6.4 Number of Relays Used: Defending against a single attacker with 1000 clients shows increased relay counts until the attacker is found.	173
6.5 Number of Relays Used: Defending against a single attacker with a smaller number of clients shows the relay count to be smaller than with 1,000 clients.	174
6.6 Percentage of Failed Transactions with the streaming attack: The attack begins at 160 seconds and is repeated with 10 malicious clients arriving every 160 seconds. A legitimate client arrives every 1.6 seconds and stays connected for 400 seconds. As time progresses, the risk profile of the legitimate clients decreases leading to a better isolation of the malicious nodes. Consequently, the PFT decreases compared to the initial burst of attack.	175
6.7 Number of Relays vs Number of Malicious Nodes: The number of legitimate clients is kept fixed at 1,000 and the number of malicious clients is increased. All the clients are present at startup. The increase in the number of malicious clients is met with automatic and progressive increases in the number of relays, to isolate the malicious nodes.	176

Figure	Page
6.8 PFT vs Number of Malicious Clients: The number of legitimate clients is kept fixed at 1,000 and the number of malicious nodes is increased. All the clients are present at startup. The increase in the number of attackers causes the average PFT over the attack simulation window to grow as a larger-sized attack impacts more client-connected relays. . . . .	177
6.9 Number of relays over time for varying CRPA factors: The number of relays utilized peaks to a much higher number with high CRPA values, however the attacker is found much more quickly. The factor controls the growth rate in the number of relays created in response to an attack. .	178
6.10 Comparison of DoSE to Epiphany and Epiphany+Speakup: At low ratio of number of attackers to legitimate nodes, DoSE outperforms other solutions, while at mid to high ratios, Epiphany+Speak-up outperforms DoSE. . . . .	179
6.11 Comparison of DoSE to MOTAG. The time it takes to find all of the insiders is significantly lower in DoSE, especially as the number of insiders increases. This is due to the smart relay management which divides the clients during successive attacks in DoSE while MOTAG relies on a stateless classification which lets intelligent adversaries fool the system.	180
6.12 Comparison of DoSE to MOTAG. DoSE is able to keep the number of failed transactions lower during the simulation because it better isolates and identifies the attacker more quickly, resulting in less clients co-located on relays with insiders. . . . .	181

## ABSTRACT

Wood, Paul Ph.D., Purdue University, August 2016. Improving the Resilience of Cyber-Physical Systems under Strategic Adversaries. Major Professor: Saurabh Bagchi.

Renewable energy resources challenge traditional energy system operations by substituting the stability and predictability of fossil fuel based generation with the unreliability and uncertainty of wind and solar power. Rising demand for green energy drives grid operators to integrate sensors, smart meters, and distributed control to compensate for this uncertainty and improve the operational efficiency of the grid. Real-time negotiations enable producers and consumers to adjust power loads during shortage periods, such as an unexpected outage or weather event, and to adapt to time-varying energy needs. While such systems improve grid performance, practical implementation challenges can derail the operation of these distributed cyber-physical systems. Network disruptions introduce instability into control feedback systems, and strategic adversaries can manipulate power markets for financial gain. This dissertation analyzes the impact of these outages and adversaries on cyber-physical systems and provides methods for improving resilience, with an emphasis on distributed energy systems.

First, a financial model of an interdependent energy market lays the groundwork for profit-oriented attacks and defenses, and a game theoretic strategy optimizes attack plans and defensive investments in energy systems with multiple independent actors. Then attacks and defenses are translated from a theoretical context to a real-time energy market via denial of service (DoS) outages and moving target defenses. Analysis on two market mechanisms shows how adversaries can disrupt market operation, destabilize negotiations, and extract profits by attacking network links and

disrupting communication. Finally, a low-cost DoS defense technique demonstrates a method that energy systems may use to defend against attacks.

# 1. INTRODUCTION

New distributed cyber-physical systems improve and optimize grid-scale energy consumption via wide-area communications. When they are attacked or disrupted, many of these systems may fail or support adversarial profits. This thesis analyzes and improves the reliability of such cyber-physical systems.

## 1.1 Motivation

Cyber-physical systems (CPS) are an emerging class of systems which integrate physical control and observation over potentially wide areas via cyber networks. These systems promise increased efficiency, interoperability, and ease-of-use of a variety of domains, including smart grid (SG), by changing the time scale of highly interactive processes such as price negotiation from hours or days to minutes or seconds. This scale allows for finer-grained control and thus efficiency improvements, especially in the domain of energy consumption. One area of CPS growth is the smart grid (SG). Consumers and producers of electric power have increasingly become distributed, especially in regions where rooftop solar panels are common place, and there is ample financial benefit to introducing fine-grained, rapid control.

The expansion of CPS is not without risk, however, as processes that were once slow and easily observed become fast and buried in complexity. Consequently, the resilience of physical systems when combined with traditional cyber vulnerabilities must be well studied. Broadly, this dissertation examines CPS resilience when faced with security threats and network faults and provides metrics and strategies for improving the dependability of such systems.

To study CPS security, the electric power domain has been selected as it grounds the dissertation in a realistic economic and technical environment. In the context of

SG, large scale CPS is used to coordinate the production and consumption of electric power, often via energy markets. These markets exist as a form of homeostatic control [1] that optimizes power consumption and generation via incentive signals (market prices).

Traditional power markets are comprised of relatively few market players. A few large generation companies and distribution utilities negotiate prices well in advance of power delivery because of physical generation constraints. Some generators respond slowly to transients because of thermal inertia, and sufficient notice must be given to allow for adequate power availability. This problem is exacerbated when renewable generation sources permeate the grid because of the inherent unpredictability of natural energy sources such as wind. Presently most demand in the grid is inflexible—prices are negotiated months in advance and regulated. For this reason, expensive sources of standby generation are scheduled to maintain grid security, ready to deliver power during system transients. New technologies and the growth of SG/CPS transition consumers from a static to dynamic market where power prices are constantly negotiated. In such a system, standby generation is replaced by demand-response (DR)—consumers provide less demand rather than generators providing more power.

The shift to dynamic markets and demand-response improves resilience and efficiency by leveraging wide-area communication networks and potentially insecure control devices to optimize load and generation. This dependence can undermine the gains in efficiency when networks are disrupted through denial of service attacks, insecure devices are taken offline, or compromised devices collude in the market. A traditional system might have a backup generator available to respond to voltage or frequency drops in the grid (independent of any communication network) and thus survive most transient events. A dynamic system, however, may rely on communications to shed loads, and a DoS attack on the market during a transient can have potentially severe consequences.

Furthermore, the explosive growth of rooftop solar in the Southwest United States has been ratcheting up demand for SG technologies. When solar and wind generation

operate in traditional power markets, the power distribution utility typically absorbs the cost of solar’s uncertainty by raising power rates in the region. Generation shortfalls (e.g. caused by clouds) place high demand on expensive on-demand energy sources. The consumer is generally shielded from the expensive energy price through regulated markets. Recently, however, utilities like SRP [2] have begun to impose penalties on homes with solar panels to recover parts of these rising costs. Smart grid technologies such as demand response and real-time pricing systems that bridge the gap between consumer loads and market prices can alleviate these penalties by optimizing energy utilization via wide area controls. If this technology is successful, it can support additional solar panel integration. The system may not be successful, however, if the cyber components of the CPS introduce additional faults and instability that roll back the efficiency gains of the power markets. The remainder of this dissertation analyzes and measures these risks and provides techniques to alleviate some of them.

## **1.2 Contributions and Outline**

This section summarizes the contributions and outline of this dissertation.

### **1.2.1 Energy CPS Models and Attack Strategies**

The first contribution in Chapter 2 is a profit-oriented model for energy-based CPS markets that supports modeling strategic adversaries. Prior work in real-time energy markets does not consider the presence of adversaries who attempt to game the market for additional profit. A new model is created that distinguishes between collective social welfare and individual profits by introducing deregulated power markets with multiple independent participants. Intertwined in this chapter is background material that explains the underlying energy system on which this dissertation lies.

When the CPS model is viewed from an individual-player profit viewpoint, an interdependent defensive game can be introduced to analyze strategies against at-

tackers. This contribution is a game theoretic approach to optimizing defensive investments in the CPS when multiple defenders exist under a strategic adversary. With limited budgets in hand, the attackers and defenders have an interest in optimizing their strategies. Interdependent defensive games, however, exhibit strong positive and negative externalities due to investments. This complexity is further exacerbated by the competitor elimination problem—the owner of an asset may be different than the most impacted party. In this contribution, an investment optimization game is constructed and examined which accounts for these properties in attack and defense. The game is further explored by introducing noise into model viewpoints and allowing defensive collaboration or pooling of resources.

With these two pieces, a background and underlying profit model is established to capture the profit motivations of strategic adversaries. Up to this point, the attacks are more theoretical in nature—they exist at the energy-flow level. The next chapter introduces technical market operations that utilize realistic communication interfaces to negotiate energy flows and prices.

### **1.2.2 Energy CPS Operation**

Chapter 3 maps the principles of market operation to realistic system implementations. The market players from Chapter 2 now exist in a market implementation that models the actual negotiation processes between energy producers and consumers. An existing market solution, dynamic market mechanism (DMM) [3], is analyzed in a real-time communication environment. The response of DMM is dependent on the state of the networking subsystem, and characteristics such as latency, jitter, and dropped messages all have a negative influence on the system’s performance. It is demonstrated that low-level network attacks can influence market pricing, and that network disruptions can potentially crash the energy market.

A second model, a novel contribution created for this dissertation, is also included in this chapter. A Nelder-Meade (NM) optimization technique is modified to support



non-stationary objective functions such as time-varying energy demands. This modification enables NM to serve as the control system driving a real-time power market. With this system, which is designed to operate on low-level distribution, can be used to analyze the profits of network attacks on energy markets.

These communications models enable attack/defense studies at the network layer. This bridges the gap between the more theoretical attack/defense models in Chapter 2 with the practical implementation in this chapter. The next chapter evaluates the combination of attack/defense strategies on practical market mechanisms.

### 1.2.3 CPS Attacks and Market Manipulations

In Chapter 4, the energy CPS operations are attacked and evaluated by two different attack strategies. In the first strategy, the attacks from Chapter 2 are applied to the Nelder-Meade based market model from Chapter 3. The attack targets in the theoretical model are mapped to DoS/outages on specific network links in the market mechanism. The profit and defense strategies are evaluated, and it is shown how the strategy can increase attacker's incentive in a realistic system implementation.

An additional, heuristic-based attack is included in Chapter 5. An adversary launches DoS attacks in an attempt to manipulate market prices while buying and selling energy from a storage device. It is shown that an adversary with access to a storage device can extract additional profits by launching such attacks. Some defensive techniques are introduced in this chapter as a method to improve resilience.

The next chapter introduces a low-cost distributed denial-of-service (DDoS) defense mechanism that can facilitate economical defense of CPS systems.

### 1.2.4 Low Cost DDoS Defense

Chapter 6 presents Denial of Service Elusion (DoSE) as an inexpensive method for mitigating network layer attacks by utilizing cloud infrastructure and content delivery networks to protect services from disruption. DoSE uses these services to create a

relay network between the client and the protected service that evades attack by selectively releasing IP address information. DoSE incorporates client reputation as a function of prior behavior to stop attackers along with a feedback controller to limit costs. We evaluate DoSE by modeling relays, clients, and attackers in an agent-based MATLAB simulator. The results show DoSE can mitigate a single-insider attack on 1,000 legitimate clients in 3.9 minutes while satisfying an average of 88.2% of requests during the attack.

This defense technique facilitates the economical defense of CPS from DoS attacks.

Chapter 7 concludes the dissertation and provide directions for future work. Additional material in Appendix A describes the simulation frameworks used throughout the dissertation.

### 1.3 Work Publication

This section covers previous, current, and planned publications supporting this dissertation.

#### **CPS – Completed Works:**

- **Interdependent Defensive Games in CPS**

Paul Wood, Saurabh Bagchi and Alefiya Hussain – Presented March 23rd, AAAI 2015 Spring Symposium, Applied Computational Game Theory, 2015

- **Optimizing Defensive Investments in Energy-Based Cyber-Physical Systems**

Paul Wood, Saurabh Bagchi and Alefiya Hussain — 20th IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems (DPDNS15), 2015

- **Defending Against Strategic Adversaries in Dynamic Pricing Markets for Smart Grids**

Paul Wood, Saurabh Bagchi, and Alefiya Hussain — 8th International Conference on Communication Systems & Networks (COMSNETS16), 2016

- **Denial of Service Elusion (DoSE): Keeping Clients Connected for Less**

Paul Wood, Christopher Gutierrez, and Saurabh Bagchi — 34th International Symposium on Reliable Distributed Systems (SRDS15), 2015

**CPS – Works Under Review:**

- **A Framework for Evaluating the Resilience of Dynamic Real-Time Market Mechanisms**

Paul Wood, Dylan Shiltz, Thomas R. Nudell, Alefiya Hussain, Anuradha M. Annaswamy — (Under Review) IEEE Transactions on Smart Grid, Submitted October, 2015

- **Attacks and Defense for Real-Time Price Signals in Smart Grids**

Paul Wood, Saurabh Bagchi, Alefiya Hussain — (Under Review) IEEE Conference on Communications and Network Security, October 2016, Submitted April, 2016

**Related Completed Works:**

- **Synthesizing and Specifying Architectures for System of Systems**

Kenley, C. R., Dannenhoffer, T. M., Wood, P. C., & DeLaurentis, D. A. (2014, July). In INCOSE International Symposium (Vol. 24, No. 1, pp. 94-107).

## **2. ENERGY-BASED CPS MODELS AND HIGH-LEVEL ATTACK STRATEGIES**

This chapter covers power flow economics, vulnerabilities, and attack/defense strategies in energy-based cyber-physical systems. The chapter first addresses the economics behind deregulated power markets as a method for identifying valuable CPS targets. A strategic adversary is then created in the latter half of the chapter to capitalize on those targets. Much of the work in this chapter was presented at the Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium in March 2015 and published in DPDNS'15 [4].

### **2.1 A Profit-Oriented Model for Energy CPS**

#### **2.1.1 Introduction**

Industrial control systems are becoming more interconnected throughout all domains and across corporations. Resources processed thousands of miles from their consumption points may traverse multiple independent companies before arriving at customers who are relying on them to operate, and these systems are gaining a large cyber footprint as automation and efficiency improvements drive a demand for internetworked components. This increased footprint creates more opportunity for malicious actors to penetrate and manipulate CPS's, especially when connected via the Internet. Successful attacks are becoming more visible [5], as demonstrated by Stuxnet [6] and shown in a recent ICS-CERT [7]. Corporations and industry needs an analysis framework and decision support tool to aid in understanding intentional attacks on interdependent CPS, their propagation through interconnected systems, and the impact they have on profitability when financially independent but intercon-

nected companies face attacks. With such a framework, defensive strategies can be formulated for improving system security and minimizing the impact of intentional attacks.

Understanding and measuring the complex interactions that occur in interdependent cyber-physical systems and creating an optimal response to attacks is a crucial step toward the goal of optimizing system profitability in the face of rising security threats. In a CPS with cross-domain interlinks, such as the natural gas pipeline and electric power generation systems, identifying high risk components and making good design choices is no longer a trivial or self-contained task. The large network of feedback created by corporate profit optimization complicates risk assessments, especially when multiple companies are competing for revenues and relying on the same input resources. Enron demonstrated in the 2000 California Power Crisis [8] that carefully placed outages can net huge profits, and understanding where and how potential attackers can profit is crucial to defending an interconnected CPS. The hypothesis is that attackers who have profit-seeking motivations will attack a different set of targets than one who seeks to simply disrupt the system.

The work presented in this section captures and models the interactions of independent CPS and analyzes the impact of cyber manipulations and induced outages on energy-based CPS components from the perspective of the overall system's and individual companies operational revenue. This lays the foundation for a utility function for subsequent sections. Using this function, experimental analysis necessitates an independent actor model, motivates profit-based objectives, and evaluates several security strategies.

To perform impact analyses, a model is created which abstracts the low-level details of components in the CPS into high level flow graphs that capture the interdependent interactions in the CPS. The graph models the high level inputs, components, and outputs that each CPS operator utilizes for profit maximization. The natural gas-electric interdependent CPS motivates a translation framework for converting perturbations in the physical system to changes in efficiency, capacity, and cost in the

graph. These parameters, with the addition of revenue and demand, formulate an optimization problem that minimizes the sum of costs, for both the individual players and separately the entire system. With this optimization, impacts are analyzed by comparing the costs of contingencies when failures or attacks cause capacities or efficiencies to change in the system.

In this chapter, the capacity, transmission costs, and revenues are modeled empirically from data collected from the United States Energy Information Administration (EIA). A model of 6 US states' natural gas and electric power systems is used to show three key results. The first is that the independence of CPS operators creates suboptimal risk assessments when performed in isolation. The second is that choosing between profit maximization and shortage minimization results in two different defensive strategies, justifying the need for economic incentive inclusion in defense optimizations. In the final result, some security strategies are considered to determine whether or not greedy attack/defense strategies based upon easily observable features of the system are optimal. The experiment shows that there is no strong correlation between two graph-observable parameters and the financial impact on the system which motivates the need for game theoretic approaches to optimizing defensive strategies.

### **2.1.2 Background**

In this section, the relevant concepts required to understand the design of the impact analysis technique are introduced. The interdependence between the natural gas and electric system are discussed along with the malicious attacker's motivations.

#### **Natural Gas System**

Natural gas (NG) is a popular fuel used in heating, lighting, electric power generation, and transportation. Its production points are often far away from load locations, so a system of transmission pipelines bridges this gap. The pipelines, which rely on

gas-fired compressors, operate as a lossy but self-contained network. Gas companies seek to optimize delivery by minimizing extraction and pipeline costs by selecting the least-loss paths. Figure 2.1 captures the high level NG infrastructure, where gas flows from production wells through a transmission and distribution system and to the end users.




Fig. 2.1. High level layout of the natural gas infrastructure.

## Electric Power System

The electric power system or grid is composed of generators, transmission and distribution components, and consumer loads. Each generation source has a different capital and operational cost which is tied to geographical features and climate, fuel type and availability, and emission control designs. A set of generators, tied in to the transmission system, responds to the instantaneous demand created by the loads and attempts to provide the lowest cost power possible based on the transmission system's capacity and each generators operational cost. For this reason, bulk power may be transmitted over long distances away from cheap generation sources toward concentrations of loads. Figure 2.2 captures the high level electric infrastructure where power is generated, transmitted, and distributed to the load.




Fig. 2.2. High level layout of the electric infrastructure.

## Natural Gas - Electric Interdependence

The electric power system relies heavily on natural gas-fired generators due to their low environmental impact and agility in adjusting output power. Nuclear and coal-fired sources often have day-long response times to transients, while NG-electric generators can be activated within minutes. Because of the cost of fuels, natural gas generators typically serve the grid's transient needs and respond to unexpected outages experienced by other generation sources while coal and nuclear sources run at a constant output. For these reasons, faults that the electric system experiences are often translated into demand spikes in the gas network. Additionally, outages in the gas network can severely hamper the electric power system's ability to respond to events that it experiences [9]. This growing interdependence creates a complex problem where impact analysis must span multiple CPS operators.

## Attacker Motives

While some attacker motives are political, the work in this chapter focuses on the financial motivations that an attacker has. Enron showed that outages can cause supply shortages and thus spikes in prices that benefit power producers on the spot-market—those selling power at on-demand rates benefit from supply loss. An easy



avenue for profit is to speculate on the futures market for power delivery. An attacker could purchase low-priced futures, cause a disruption, and then sell them in the starved market for a profit. Another alternative is to short-sell stocks from public power utilities and subject them to some outages to reduce profits and lower stock values. While the executives of Enron faced federal prosecution, cyber attackers and cyber forensics face more geopolitical barriers to be a proper deterrent of attacks. Therefore, companies must make smart investments to defend their CPS's.

CPS designers must consider security from two perspectives. The first is economic incentives (EI) and the second is resiliency control (RC). The important distinction between the two categories is the method by which a defender needs to predict failure. An uninformed RC failure may occur randomly, based on a probability distribution driven by physics and the environment, however an economically incentivized attack will focus specifically on the most financially impactful components in the system. This motivates the creation of a utility function that has a basis in revenues and costs as opposed to an approach which only evaluates shortages or random failures. Since resiliency control impacts can be forecasted, they are not the focus of this chapter.

## **System Security**

The security of CPS revolves around a set of defensive and resiliency investments. The core defenses are techniques such as encryption, intrusion detection, and malware detection while core resiliency techniques focus on redundant, fault tolerant equipment and spare capacity. Once a security strategy is developed, with the support of an impact function, system architects can incorporate a realistic risk when deciding which technologies to deploy and to what systems.

### 2.1.3 Related Work

#### Utility System Modeling and Security

Minimizing costs in electric and gas utilities is not a new problem, and several solutions have been created and utilized to solve these optimizations. They are not suitable, however, for multi-player games. These techniques model each utility system and produce device-level operational scheduling (unit commitment) aimed at minimizing cost by using varying optimization techniques such as mixed integer linear programming (MILP). When network conditions such as a minimum power delivery rate are added to the unit commitment model, a security constrained unit commitment (SCUC) is created [10] which improves system operation during failures. These optimization problems are useful for planning the day-to-day operation of each utility, and the SCUC has been extended to include a natural gas-electric combined optimization [11–14]. While these techniques produce an optimal unit schedule, they do not optimize the multi-player objective function as done in this dissertation. Instead they consider the system as a monolithic (single player) entity for optimization purposes. Market processes such as those utilized by New York Independent System Operator (NYISO) allow power producers to bid into power markets. The model does not consider profit, however, as the system attempts only to minimize cost not maximize profit.

#### Graphical Approaches to Security

Graphs have been used extensively in solving cyber security problems, but most approaches do not consider the continuous attacker objectives seen in attacks on the power utility CPS. Most of this related work focuses on protecting a particular component of a system where an attacker either has control or does not [15, 16] and identifying at-risk components for protection using this assumption. Graph-based metrics are often applied to power system problems [17], but the utility in identifying

vulnerable targets by using static vertex or edge properties may be misleading [18]. In this work, dynamic and interdependent properties are considered in attack and defense (i.e. actual impact instead of estimated). While the existing approaches are useful for modeling the defense of particular components or observing how an attack may progress across different cyber components [19], the problem of developing a security strategy has not been adequately addressed for interdependent CPS.

### **Game Theory in CPS**

Once an impact understanding is developed or assumed, game theory approaches can be utilized to improve defensive decisions. Several recent techniques [20–24] have been evaluated for defense application in the power grid. These approaches provide a good framework for using game theory to solve problems in CPS security, and this work aims to improve on the utility functions available. Little work has been done in utilizing game theoretic approaches to evaluate the combined interdependencies of gas and electric networks, however. More research is needed to understand how to best defend interdependent CPS networks against an attacker especially when the defensive objectives of each player is different.

Interdependent security games, as surveyed in [25, 26], optimize and study the process by which defenders invest resources to mitigate and minimize risk from cyber attacks. This class of games involves multiple defenders in the face of attacks, excluding single defender-attacker games, and focus on the defensive investment strategy [27, 28]. Specific games are targeted toward networked control systems [29], however these models focus on the contagion [26, 30] aspect of security rather than the market-level interdependence that this model considers.

The financial impacts of attacks have been studied in [30]. A model is created which has the property that defensive investments generally reduce the ability of the attacker to sustain attacks. However, the interdependence effect only modeled

positive externalities and does not include the negative externalities modeled in this section.

#### **2.1.4 Overview**

The section provides an overview of the interacting forces in interdependent CPS's and how they will be captured in a model. The actors in the system are defined along with their independent objectives, and the threat model and device level impacts are outlined.

#### **The Actors and Objectives**

The interdependent CPS in the natural gas-electric scenario is comprised of several gas and electric entities. Gas is extracted by drilling companies, transported by pipeline operations, and distributed to customers, and similar actors exist on the electric side. Each group of actors has its own customer base, which may be another actor in the system or direct consumers, and tries to maximize its profits as a cost optimization objective.

#### **Actor's Resources and System Model**

Each actor owns a set of components that are interconnected with the larger CPS. These components perform one of four functions, and the first is production which introduces resources from external, unmodeled sources into the CPS. The second is consumption of a resource in the system, which acts as load. The third component is the transportation or transmission function such as a power line or gas pipeline. The final function is a transformational operation which converts a resource from one form to another. These components are interconnected and divided among actors to create a model of the CPS.

## Threat Model and Impact Analysis

The attack motivations categorized as economic incentives (EI) and resiliency control (RC) drive distinct classes of failures. Under the RC umbrella, mechanical degradation, random occurrences, and unplanned events act to perturb the actor's resources. A NG pipeline compressor may fail due to a worn bearing for example. Under the EI umbrella, physical and cyber attacks actively seek to disrupt hardware in the system with a specific agenda. The occurrence of failures under EI is systematic, informed, and potentially widespread. In both cases, however, the actual mechanics of an attack or failure translate into parameters in the system model. An equipment piece that fails causes reduced capacity or increased loss in the system which maps to suboptimal flows and reduced profits.

Given the system model and no perturbation, the most profitable flows can be established by the actors. An attack can then reduce the capacity or increase the loss of a component in the system model, resulting in a new profit. The impact analysis then is the difference in profits for these two models.

### 2.1.5 Approach

In this section, the system model and optimization problem is formally described. The resources are normalized to allow inter-domain comparison and evaluation. A graph structure is defined that maps the core parameters of each component in the system to a graph vertex or edge parameter. An optimization problem is then formulated for solving single and multi-objective system models based upon this formulation.

## Nodes

Nodes or vertices in the graph serve either as hubs, sources, or sinks, and these components provide the interface with non-modeled processes.

**Hubs** The hubs in the graph represent zero-sum routing components in the system such as electric buses or gas distribution headers. The hubs capture points in the system where alternate transmission paths or sources can be selected and enable the splitting and combining of flows. Hubs also serve as geographic anchors and act as abstraction points where detail is removed from the underlying system. A distribution system which has both load and generation can be modeled as a hub with an in and out flow to capture all of the low level electric bus impacts, for example. The hub's geographic property is used to calculate distance-based transmission losses.

**Sources and Sinks** Some nodes act as sink or source points in the system, and these nodes are the interface between components not in the model. The various sources of energy (gas wells, flowing water, coal mines) are the input sources to the model and represent components that have a capacity and fixed unit flow cost. The sinks in the system are loads or consumers of energy, and these nodes absorb flow and generate revenue for the system and it is where the model terminates.

## Edges

The links between most nodes model resource transportation components such as gas pipelines or electric transmission lines. These edges have some capacity and resource loss due to inefficiencies such as resistance. In the case of an electric transmission, the line losses result in less energy at the output than was provided at the input as a function of flow. Similarly, gas loses pressure in transmission as it flows. Thus each edge has four associated parameters: capacity  $c(u, v)$ , flow  $f(u, v)$ , cost per unit flow  $a(u, v)$ , and loss  $l(u, v)$ . Flow is the output of the optimization, and cost is associated with edges connecting sources and sinks. For most edges, there is only a capacity and loss parameter.

**Electric Transmission Model** Each high voltage transmission line has a designed capacity that maps to  $c(u, v)$ . The transmission line losses are a function of line length,

voltage, and electric current flow and the magnitude of loss can be estimated [31] as a ratio of percent per distance. A transmission edge connecting node  $u$  and  $v$  has a loss per unit flow defined as  $l(u, v) = \text{lossrate}(u, v) \times \text{dist}(u, v)$ .

**Fluid Transmission Model** Similarly, by Bernoulli’s principle, the pressure in a gas pipeline will be reduced as it flows requiring compressor stations to boost pressure. As in the power case, this cost can be calculated as a distance dependent loss. The capacity of the pipeline is also a known design parameter and is planned based on long term demand forecasts made by the gas utilities.

**Transformation Edges** The remaining transformation edges convert the energy in the CPS from one type to another. This transformation is identical to transmission with the caveat that efficiencies are much lower. In the gas-electric example, a transformation edge converts natural gas into electric power with typical efficiencies below 65%.  $l(u, v) = \text{Conversion Efficiency}$ .

## Cost, Revenues, and Expenses

Certain links in the system have specific costs or revenues associated with them. The operational costs for each type of power generation plant contributes to the expenses while the consumers of energy provide revenue. These money flows are captured by a positive or negative value of cost  $a(u, v)$ . An additional function, demand  $d(v)$ , dictates how much a energy a consumer is willing to purchase or supplier is willing to sell. The set  $L$  (loads) contains all of the vertices  $v$  for which  $d(v) > 0$ , and are sinks. Additionally each source has the property  $d(v) < 0$  and comprises set  $G$  (generators).

**Resource Normalization** Resource unit conversions are required to normalize pricing information based on different quantities—gas is often priced in a volume while electricity is priced per joule or watt-hour. A system-wide view requires a normalized

unit of measure for all of the components in the system. In the case of energy based systems, i.e. power and gas utilities, the standard SI unit of joule is be used and converted to megawatts (MW), but other units can be used for other systems.

### Utility Function

The components are combined into a flow graph representation, as shown in Figure 2.3. Linear programming and minimum-cost flow algorithms are then used to analyze the system. The utility function for the entire system is defined as follows. Each edge in the graph has a capacity  $c(u, v)$  and cost  $a(u, v)$  defined by the owning actor as a property of the physical system. The vertices  $V$  contain all of the distribution headers or power buses, plus any sink or source nodes. The following constraints are established and then linear programming is used to optimize the flow solution.

$$\text{Utility} = \min \sum_{(u,v) \in E} a(u, v) \cdot f(u, v) \quad (2.1)$$

Subject to constraints:

$$0 \leq f(u, v) \leq c(u, v) \quad (2.2)$$

$$d(v) \leq \sum_{u \in V} c(u, v) \text{ for all } v \in L \quad (2.3)$$

$$s(v) \geq \sum_{u \in V} c(v, u) \text{ for all } v \in G \quad (2.4)$$

$$\sum_{u \in V} f(u, v) \leq d(v) \text{ for all } v \in L \quad (2.5)$$

$$\sum_{v \in V} f(u, v) \leq s(u) \text{ for all } u \in G \quad (2.6)$$

$$\sum_{w \in V} \frac{f(u, w)}{1 - l(u, w)} = \sum_{w \in V} f(w, u) \quad \forall u \quad (2.7)$$

The first equation 2.1 measures the cost across all edges in the system, and its minimization optimizes system-wide profits. Equations 2.3 and 2.4 constrain the






Fig. 2.3. Flow graph of the combined natural gas and electric infrastructure.

potential demand values to ones that satisfy the input or output capacity of the node, and equation 2.2 ensures the flow solutions are less than the edge capacity. Equation 2.5 and 2.6 prevent over-selling to a load or over-production from a source. Equation 2.7 is the conservation of energy at the intermediate hubs. The division by  $1 - l(u, w)$  on the left hand term captures the impact of transmission losses. The sum of the inputs to a node will be larger than the outputs to account for the fact that some of the input energy is lost in transit to the hub. The net cost of these losses is dependent on the source cost, and placing the losses here ensures the optimal flow of energy and causes the hubs to have a non-zero sum.

## Utility for Multiple Actors

The notion of multiple actors arises as a function of asset ownership and financially competing entities. Each asset (edge) is owned by an actor (each edge has one owner), and these owners are acting in a competitive market. The actors are autonomous entities and will not cooperate or form cartels for profit maximization. Instead, a cartel of individual actors will be modeled as a single actor.

Estimating the utilities in the presence of multiple actors entails a large body of economic works that estimate supply and demand pricing in open markets. The complexities of such assessments are not necessary for an accurate model of utility provided certain assumptions are made. The first assumption is that if the graph structure remains the same, then the addition of new actors or changes in ownership do not change the optimal flows in the system. The rationale is that, given a ground truth set of costs and losses, a competitive i.e. non-collusive set of actors will eventually settle on the most cost-effective flow, since the lowest cost option provides maximal profit system-wide (the system will enter a coalition-proof Nash equilibrium). The second, derivative assumption is that since the system's sink and source flows and prices are fixed, the net system profit is independent of asset ownership. The sum of the individual actors profits equals the original system profit after any changes in ownership. The rationale is that while some profit distributions may not be entirely credible, it relaxes profit distribution from a negotiation problem to an assignment problem. The final assumption is that costs in the system include required minimum profits for operation.

The problem for determining the utility in a multiple actor scenario is now viewed as a fair profit assignment algorithm. The fairness piece comes into play by observing what an actor can charge at each point in the system. Notionally this is done by raising prices until the buyer goes somewhere else. Practically this is measured by constricting the flow out of each actor, independently, and observing how much more the system as a whole is paying to supplement that reduction (the system's marginal

cost). Since the prices at the end suppliers and consumers is fixed, and the flow through the system remains unchanged due to the allocation of internal costs, the marginal cost can be applied as a profit allocation. For example, if an actor constricts her flow on a particular asset from 50 units to 49 and observes the system's marginal cost as \$1, then he determines he can charge the original price + \$1. In the model view, the cost  $a(u, v)$  on the edge goes from \$0 to \$1 while the system flows are unperturbed. The system's net profit is decreased because this is a cost, however, the actor's individual profit is increased by that amount. The result is a profit distribution that follows the assumptions above.

The marginal cost alone, however, cannot be used to distribute profits fairly because a reseller situation may arise. Imagine three independent actors are operating in a series. Each actor determines its marginal cost to be \$1. This cost was calculated, however, with the assumption that the other actors in the series would take zero profits. Intuitively, the profits available at the last actor in the series must be split among all three so that the total system profit remains constant. The fairness in this situation is taken by attempting to uniformly distribute the profits among actors in these situations. The solution approach, as listed in the proceeding algorithm, is to grow from small to large fractions of the marginal cost uniformly across the actors.

**Limiting Information** Certain situations may arise where the model needs to be explored from the perspective of an actor or adversary who has a restricted view of the system. This is the case when someone estimates capacity for example. The application of this knowledge level is to add noise to the different parameters in the system, with a feasibility restrictions. The simplest way to do this is to center a normal distribution around the mean of the original graph parameter and then vary the standard deviation based upon the knowledge level.

For each parameter, except for the flow  $f(u, v)$ , the knowledge level  $\sigma$  is defined as  $c'(u, v) = \mathcal{N}(c(u, v), \sigma^2)$ . The distribution is truncated such that the signs of the




Fig. 2.4. The graph shows an example scenario where each asset is owned by a different actor A-F. The three suppliers A, B, D have unit cost of production \$3, \$3, and \$9 respectively and the end customer has a fixed unit price of \$11. At each edge, there is a fixed unit cost and a marginal unit cost. The fixed cost is a parameter of the model, while the marginal cost is calculated by observing the system's response to constricting flow through that particular edge. In this case, actors A and B are in direct competition, so no unilateral price movement is possible. C, however, is in a position to mark up the price an additional \$6 because it is in competition with D. E and F are in series and may mark up to the remaining purchase price at the customer, cumulatively. A fair split of \$1 for E and \$1 for F is shown as an example.

original parameters do not change. The new flows are determined by evaluating the optimization problem for the new parameter set.

## Measuring Impact

Each parameter, cost  $a$ , capacity  $c$ , and loss  $l$  can be varied to determine the new utility.  $\text{Impact} = \text{Utility}' - \text{Utility}$  where  $\text{Utility}' = \text{Utility}$  as  $a, c, l \rightarrow a', c', l'$ . The impact is computed when an attack or other natural forces act on the systems. The impact can be negative to indicate a gain in the Utility when there are multiple actors in the system. In the models, the gain is primarily due to actors selling spare capacity in the event of a failure.


### 2.1.6 Experimental Model

#### Geographic Model

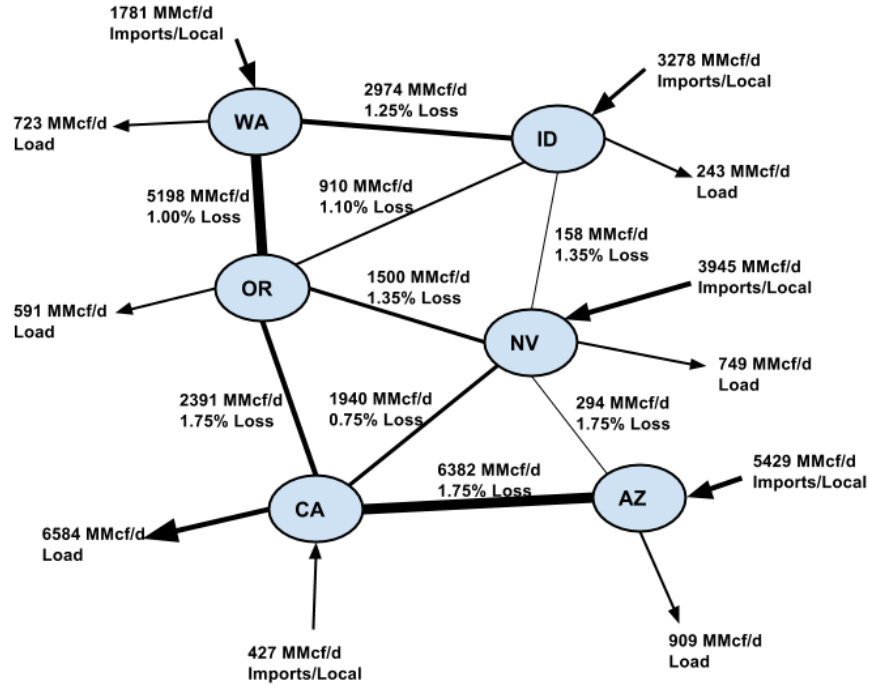
Six Western US states are captured in a moderately complex, interdependent gas-electric system model. The state-level representation maps to data available from the Energy Information Administration (EIA) [32, 33], and the states of Washington, Idaho, Oregon, California, Nevada, and Arizona are chosen because this region is generally a net importer of energy. Each state's hub is mapped to its geographic centroid, and edges are created based on border adjacency. Each state has two energy hubs, one for natural gas and one for electricity, and a customer base for each resource. In total there are 12 hubs and 18 hub-to-hub transmission lines. Figure 2.5 depicts the infrastructure for the two systems, and the interconnection occurs between the load side of gas (b) and the generation side of electricity (a).

#### Interconnected Infrastructure Model

Four functions must be defined for hubs and edges in the gas infrastructure. The cost function  $a$  is based on the average price paid in each state over a year. For import edges, where gas is purchased out-of-model, the cost is taken to be 25% lower than the price customers pay, motivating transportation. Next the loss function  $l$  is defined. For this step, a calculation is made based on a typical loss of 1% per 400 km [34],



(a) Electric Model



(b) Natural Gas Model

Fig. 2.5. A flow model is created for six Western US states for both an electric (a) and natural gas (b) infrastructure.

since the actual loss rates vary based on each individual pipeline’s construction characteristics. The resulting loss rates are seen in Figure 2.5 (b). The capacity function  $c$  directly maps to EIA’s dataset [32]. Each state’s energy information profile details energy produced by gas-fired generators and their efficiency,  $l$ . Finally the supply, imports, and demand for each state were calculated by averaging yearly consumption into short term amounts.

Similarly, the values for the electric infrastructure are calculated using the EIA sources [33]. Each state has a suite of electric energy sources to choose from, nuclear, coal, natural gas, solar, etc., and each source has its own edge into the hub. The prices for these different sources are estimated and the supply and consumer pricing in the system is assumed static because most contracts are negotiated for terms of a day or longer [35].

## Model Adjustments

To represent a more challenging model, several modifications are done to the calculated infrastructure. The installed electric capacity  $c$  is reduced by 25% to account for inoperable generators due to maintenance and climate, and the demand is increased by 65% from the daily average to represent a high-demand period, i.e. in the peak of winter. With these adjustments, the system has about 15% spare capacity which is in line with the EIA’s spare-capacity estimates.

## Operational Complexity

To evaluate whether or not the model is sufficiently complex, the system’s response to capacity reductions is plotted in Figure 2.6. Four example edges are selected and their capacity  $c$  is reduced between 0 and 100%. The plot shows the marginal increase in cost per additional percentage of reduction. The stepping behavior represents the different alternative sources that replace the eliminated capacity. The solid line, for example, has a cheap alternative in the 5-15% reduction range. However, after 18%

the system selects increasingly costly alternatives to augment the reduced capacity. Since the lines have a variety of steps, the 12-hub model is complex enough to study.

The model created in this chapter is evaluated for a combined natural gas and electric power distribution system. Three scenarios are examined to answer several questions. The first scenario looks to establish the need for cooperation to achieve optimal defensive strategies. The second scenario investigates the difference between objectives of profit versus shortage to show that the choice in objective is crucial when analyzing an attacker/defender strategy. The final experiment investigates greedy attacker strategies to see whether or not a game framework is necessary to answer attacker strategy questions.

## **Quantitative Data and Model Creation**

The goal of these experiments is to create a realistic scenario in which to test hypotheses about the utility of different components in the system. The high level flow amounts and capacities should reflect what is seen in reality to avoid creating a problem and solution from a scenario that might never exist.

Information from the EIA was used to create a model of the gas and electric systems of the western region of the United States (California, Nevada, Arizona, Oregon, Idaho, and Washington states). The geographic centroid of each state was used as a node in the graph, and each adjacent state's interconnections were summarized into single edges. The six state region was selected because of its relative isolation geographically and from the natural gas infrastructure perspective. The region imports most of its natural gas from fields adjacent to the modeled states, and several of the states have no external edges such as Oregon and California, having only minor interactions with imports and exports of gas and electricity from outside of the system. The 12-node system (6 for each resource type) is still complex enough to capture the interactions seen in large CPS without being too large to burden experimentation.



The model uses two actors, one for the natural gas infrastructure and one for the electric infrastructure. Each actor attempts to maximize its own profit in the scenarios tested, unless otherwise mentioned.

**Natural Gas Infrastructure** The interstate pipeline capacities for each state are available in the EIA database which provides several types of state-level data. Each state has an exact import/export capacity to other states and a per-state consumption volume. This provides the capacity and consumer information for each gas node. A portion of the gas moved is utilized by electric power generating facilities, and the interconnection is captured by the transformational edges in the graph.

The next component is loss due to gas transport. These losses are based on estimates formed from the evaluation of various US FERC reports which provide information about how much each pipeline is charging for losses and uses of natural gas due to transportation. Finally, revenue information was collected from the EIA to provide values for the general gas customers versus the electrical generation customers. The electric utilities generally paid less for gas, likely due to the large contract size and optimization or predictability of demand.

**Electric Infrastructure** Similar to the gas infrastructure, the information provided by the EIA was used to construct the same set of graph components for the electric actor. The primary difference is that transmission infrastructure information is not publicly available or easily accessible, so the capacities and losses are less exact. The generation losses are based on the actual performance data of units in each state. The transmission losses are estimated based on distance. These costs are operational only and do not include the amortized capital costs associated with plant construction and lifetime because these are considered sunk costs when making daily operational decisions.

**Combined Infrastructure** The two subsystems rely on different energy measurement units for basic transport and pricing information. Since both systems deal with

energy, the standard unit Watt was used to convert all of the components into the same unit system. Gas is transported in a volume and may have different amounts of energy per cubic foot, but the variance observed in the measurements for each system revealed less than a 1% difference in energy content between sources so a standard value was taken.

To establish demand, the year average consumption of both gas and electricity is provided by EIA for each state. This amount was taken as the mean demand and increased by 20% for the gas system and 65% for the electric system to reflect the variance seen in average daily load versus peak load. Additionally the capacity used for the electric system is provided as installed capacity, not available capacity. This amount is reduced by 25% to reflect typical plant outages or reduced water supply availability.

When the system is combined, it is important to understand how the interdependencies manifest in the real system, as captured by the model. Figure 2.6 shows how the capacity reduction of four different edges in the system have substantially different impact profiles. An important concept in the power system is understanding which source the next watt of demand will be allocated from. The figure shows the marginal costs of providing the alternative watt when it is lost from the attacked edge. Low values mean that a cheap alternative exists while high values means that the most expensive fuels are utilized to power the system.

### 2.1.7 Results

In this section, three experiments are conducted to study the impact of attacks in a multi-actor interdependent gas and electric CP system. The first experiment shows that co-operative strategies across the gas and electric CPS reduce the impact of most attacks. The second experiment investigates the difference between shortage and profit impact, and the last experiment shows that greedy security strategies are not sufficient for impact analysis in interconnected CPS.

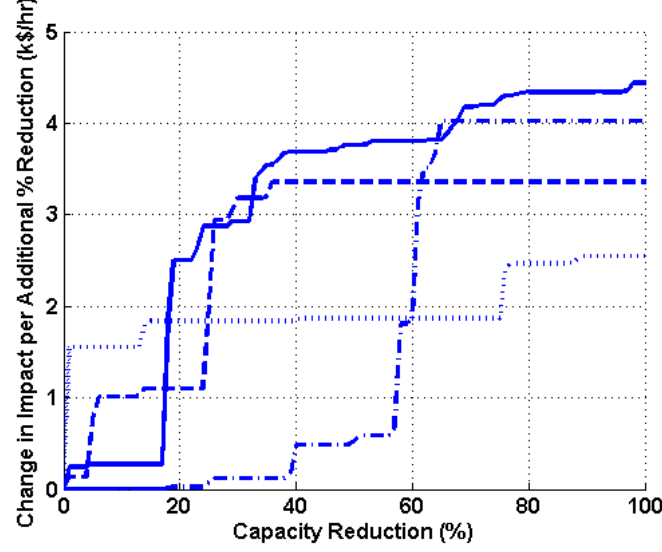


Fig. 2.6. This chart shows the marginal change in the impact for incremental reductions in capacity for four edges in the system model. The incremental steps represent the costs of the marginal generation that must be utilized to deliver energy to the customer.

### Single Failure Impact Analysis

The impact of an attack on the interconnected gas and electric infrastructure is studied. For each edge in the graph, the capacity  $c(u, v)$  of the edge is reduced to zero to simulate an attack. The impact  $I$  of the edge loss is computed by computing  $\text{Utility}'$  (Equation 2.1) and then subtracting it from the computed Utility when the edge is present.

Two scenarios are considered. In the first scenario,  $I$  is computed in the presence of an edge failure with a single actor representing the interests across the complete interconnected gas and electric network. This scenario indicates how information sharing results in lower impact and enables a more secure interconnected CPS. The impact computed in this scenario is indicated as *combined* in Figure 2.7. In the second scenario,  $I$  is computed independently, with one actor representing the gas network and one actor representing the power network. The impact computed for each actor in this scenario is then summed together and indicated as *independent* in Figure 2.7.

The following observations are made from Figure 2.7. (a) The impact of independently operating actors is higher or same as the combined gas and electric CPS. While only the results for the top 10, out of 38, edges are shown, this is also true for the other transmission and transformation edges in the model. (b) The impact is much higher for a small set of edges and then rapidly reduces for the other edges indicating that some edges are more critical than others. For example, when Edge #1 fails, which is a transmission edge to indicate gas imports into Arizona, the highest impact is observed in the combined actor model but it is lower than the independent actor model. The difference between the impact in the two scenarios, is due to redirecting spare gas resources to satisfy the NG fired electricity generators in the combined actor model where as in the independent model, the independent electric network must rely on more expensive sources of power. (c) Failure of some edges have the same impact in both scenarios. For example, when Edge #2 fails, which is a transformation edge from a gas to electric hub in California, the impact is same.

Hence this section empirically shows how a combined CPS can be more secure when the system is cooperative across different sectors optimizing resources across the full interconnected gas and electric CPS rather than considering each system in isolation.

### **Single Failure Shortage Analysis**

In this experiment, the difference between shortage experiences within the system and financial impact captured by the utility function is investigated. Based on the methodology developed in the previous experiment, the capacity of an edge is reduced to zero, and the impact is computed based on a single combined actor. Additionally, for each failure, the shortfall (that is demand not met) in watts is computed at the loads in the system. The results are shown in Figure 2.8 where the y-axis indicates the normalized financial impact and the normalized shortage impact. The normalized impact is computed by dividing the impact of the current edge failure with the

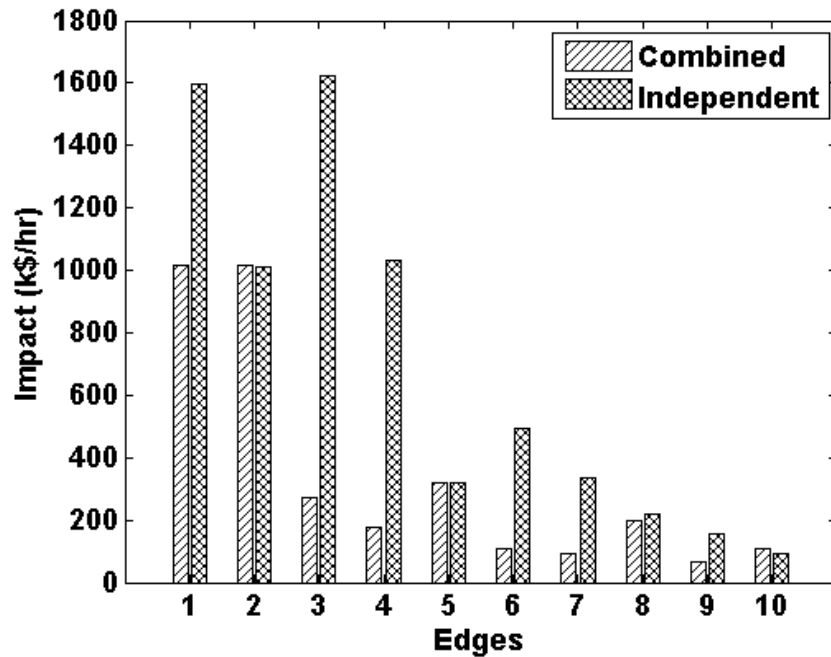


Fig. 2.7. This chart shows the impact of objective independence in ten edges due to a failure in the model. **Combined** indicates a single actor model where as **Independent** indicates two actors acting independently.

summation of the impact caused by the failure of each of the other edges. The financial and shortage impacts behave differently. For example, while financial impact for Edges #3, #4 and #5 is similar, the shortage impact is substantially different, due to the availability and cost of alternative generation resources. This indicates the need for detailed models when incorporating financial impact of shortages, by associating nuanced models to capture unmet demand in dollars. Shortages should be carefully modeled to ensure a secure CPS. If shortages have associated penalties which are not subject to a "force majeure" clause, the financial ramifications should be included in the impact analyses.




Fig. 2.8. This chart compares the normalized financial impact (left) to the shortage (right) caused during a complete failure of an edge in the system. The impact analysis could be different based on which objective is chosen, profit maximization or shortage minimization.

### 2.1.8 Security Strategy Analysis

In this section two greedy security strategies are evaluated. The first strategy is to attack the largest capacity edges first and the second is to attack the highest flow edges. If a greedy strategy is successful then a complex security strategy may be unnecessary. It is shown that for interconnected CPS, a greedy strategy is not successful since there is no monotonically increasing trend between the greedy selection metric (capacity or flow) and the impact that it has on the system.

To evaluate each strategy, a failure on one edge is modeled and the impact the failure has against the greedy selection metric is used for comparison.

Figure 2.9 compares the edge's failure impact as compared to the capacity of the edge. For this strategy to be successful, the impact should increase as the capacity

increase from left to right in the plot. However, it is observed this this not the case indicating that a greedy approach does not work in an interdependent CPS. It is also observed that there are several high capacity edges that upon failure that do not have a high impact on the interconnected CPS.

To understand the dynamics of such failures, part (b) shows the impact versus capacity reduction for two representative edges. Both there edges have similar impact but have different capacities. When a failure occurs, the interconnected CPS re-optimizes its flow with the a new least-cost plan. The impact of the alternative flow is shown as the points along the line, and its slope is driven by the cost at the source of the new resource. When the edge's capacity is reduced between 0-100%, source square may be used at a low cost, thus a low slope. Between 80% and 100%, source circle may be used which has a very high cost.

Next similar analysis with flow is conducted. The largest edges are not always completely utilized hence flow can be less than capacity. Figure 2.10 indicates the impact of the failure in edges based on the observed flow. Failures in edges causes the volume of resources to be rerouted, but because multiple redundant paths and low-cost alternatives may exist, the selection of the highest-flow edge is not universally the best. To understand the dynamics of such failures, part (b) shows the impact versus flow reduction for two representative edges. Both there edges have similar impact but have different flow. The capacity of the edge is progressively reduced to indicate a reduction in flow. The circle point has some spare capacity, as its initial flat portion in the line. The small flow redirection in the square point results in a high, steep penalty when compared to the shallow growth of the first line. The slow growth is again due to cheap alternative sources.

### 2.1.9 Conclusion

In this section, an approach to impact analysis was presented to support game theory application to a combined gas-electric utility system. The impacts of equipment




Fig. 2.9. Impact of edge failure compared with its transport capacity. A lack of a clear increasing trend suggests that this approach is non-optimal in attacking. (b) shows how the impact of the two points progresses differently as capacity is reduced. Since the impact is low for a large capacity loss in the circle point, based on the slope of the line, the attack has less overall impact on the system.

outages and capacity reductions on company profits and system-wide profits were calculated. When companies act independently, it was shown that some sub-optimal choices are made in defensive resource allocation which supports a cooperative game approach to developing security strategies. Additionally it was shown that depending on the attacker's objective of profit or disruption, different defensive maneuvers would be made, so a careful bridge between shortage-based and financial-based adversaries should be created. The game theory approach to defense was motivated by showing that naive security strategies are insufficient in determining the most valuable target, implying a more complex decision process should be involved in optimization. The security analysis of the interdependent gas-electric system is ripe for further game theory development and provide an approach for measuring the utility of different system components.






Fig. 2.10. Impact of edge failure compared with its optimized flow. The high flow edges, when attacked, result in the most redirection of energy in the system. (b) shows how the alternative energy cost may be much higher per watt for some sources than others, the greedy approach will not work as show in (a).

## 2.2 High-Level Adversarial Strategies

### 2.2.1 Introduction

This section addresses the problem of asset protection in the face of strategic adversaries in an interconnected CPS. The model that is developed in this chapter is of autonomous organizations (equivalently, corporations) dubbed as “actors”. The actors own and operate various assets, and cooperate to provide some end-user visible service. For example, the natural gas provider(s) and solar energy provider(s) feeding into the electric grid, provide electric power to end consumers. Attacks against these assets impact the profits of the actors. Motivated by the prospect of financial losses, defensive investments are made by the actors. The gamut of relationships that can exist between the actors as they relate to the defensive strategies that they deploy are explored. The gamut runs from actors behaving completely independently through a subset of them cooperating in securing the assets to perfect cooperation.

There are a few insights into improving the models for defensive investment optimization. First, the implications of attacks in the cyber side should be measured on the physical side. This enables dependencies to be drawn from complicated interconnections rather than approximated via contagion. Second, when every actor in the system is considered to be financially motivated, then attacks are driven by profits and defenses are driven by losses. This allows for adversaries to be profit seeking and creates a complication for defenders where the assets which cause the most harm to one actor may be owned by another. When actors are mutually harmed by an attack, they may wish to collaborate in defense and share the expense of defending an asset.

The solution captures the physical interconnections as a directed flow graph. The nodes and edges capture the primary supply chain factors involved in a system such as the interconnected natural gas pipeline and electric grids. These factors are the maximum capacity, cost per unit flow, and loss due to inefficiency. The flow is then optimized under a multi-actor model which measures the profitability of each actor. This model then serves as the basis for impact analysis—the supply chain factors are perturbed during cyber-attacks and the change in profitability is measured. The strategic adversary model then optimally selects a subset of actors in the system and targets which have a large positive benefit to the attacker. The defenders, estimating the adversary strategy, independently select assets to defend.

The model is evaluated against an interconnected natural-gas, electric system which is created from data available from the Energy Information Administration (EIA). The impact of multiple stakeholders is evaluated in the impact model, showing that the inclusion of independent actors significantly influences the observed impacts of cyber-attacks on asset owners. The strategic adversary model is evaluated against varying number of actors and noise to capture the adversary’s sensitivity to accurate models. Finally, the defense strategy is analyzed in its effectiveness at protecting against the strategic adversary.

### 2.2.2 Model Overview

The section provides an overview of the interacting forces in interdependent CPS's and how they will be captured in a model. The actors in the system are defined along with their independent objectives, and the threat model and device level impacts are outlined.

#### The Overall System

The desired outcome driving this work is an optimal defensive investment strategy for each actor in an interdependent CPS. The first component is an impact analysis tool which measures the financial outcomes of perturbations in the physical system, which are driven by cyber attacks. The impact analysis is then used to drive a strategic adversary who evaluates the best targets to attack, given the particular impact model. The final piece is the defender who takes the preceding two pieces and combines them to estimate an attacker's moves and counter them with defensive investment at crucial locations in the system.

Refer to Section 2.1.5 for the impact model details.

#### The Actors and Objectives

The interdependent CPS in the natural gas-electric critical infrastructure scenario is comprised of several gas and electric entities. Gas is extracted by drilling companies, transported by pipeline operations, and distributed to customers, and similar actors exist on the electric side. Each group of actors has its own customer base, which may be another actor in the system or direct consumers, and tries to maximize its profits as a cost optimization objective.

Each actor owns a set of components (assets) that are modeled in a directed graph. The graph's edges have weights which correspond to the costs and losses observed

by the system during operation, and an optimization problem is formulated that maximizes the profitability of the system.

### **Threat Model and Impact Analysis**

The threat model captures cyber attacks as actively seeking to disrupt hardware in the system with a specific agenda. The occurrence of failures are systematic, informed, and potentially widespread. The mechanics of an attack, whether it attacks a particular programmable logic controller or some other sensor, are translated into perturbations in the system model. An equipment piece that fails causes reduced capacity or increased loss in the system which maps to suboptimal flows and reduced profits.

Given the system model and no perturbation, the most profitable flows can be established by the actors. An attack can then reduce the capacity or increase the loss of a component in the system model. As a result, the system operates at a lower efficiency, having higher costs or increased losses. The impact analysis then is the loss in profits from these perturbations.

#### **2.2.3 Attacker Strategy**

In our model, the adversary takes on the role of a subset of the actors in the sense that it tries to maximize the profit of some of the actors through a cyber-attack. The rationale is that the adversary can get a share of the profits of the subset of actors, e.g., by getting equity in the companies represented by the benefiting actors or otherwise engaging in the market. Conceptually (and simplifying somewhat), one can think of an impact matrix with the actors as the rows and the assets as the columns. The value  $IM[a, t]$  implies the impact of taking down target  $t$  on actor  $a$ . If this is a positive value, it implies actor  $a$  makes a profit out of this perturbation; if it is negative, then it suffers a loss.

## Attacker Incentives and Constraints

In the impact model in Section 2.1.5, a few observations can be made which motivate the existence of a profit-seeking strategic adversary. The first is that disruptions in the system cause a re-routing of flows around the problem area, and these re-routings may shift profits from one actor to another, i.e. competitor elimination.

Assets in the impact model have several parameters that are properties of the particular design and implementation of the CPS. These parameters are the cost of attack, the cost of defense, and the probability of successful attack which are in addition to parameters in the impact model. The cost of attack represents the manpower and research required to disable an asset, the cost of defense represents the same for mitigating attacks, and the probability of successful attack plays into the expected return on investment.

## Strategic Adversary

The strategic adversary (SA) assumes the role of an actor or actors in the system and attempts to make investments with positive returns. Unlike in traditional impact models, where the SA is viewed as only seeking damage, the SA here is causing damages only to the extent that they support profits for some actors in the system. This creates a scenario where attacks to some targets, which may be very damaging, are unlikely because they do not create an opportunity for profit in the system.

## Target Selection Algorithm

Each target  $t \in T$  has an expected cost of attack  $C_{atk}(t)$  and an impact at actor  $a$ ,  $I(a, t)$ , which is positive to represent gains, and a probability of successful attack  $P_s(t)$ . The attacker's target set is  $T(i)$ , actor set  $A(i)$ , and is limited to spending  $M_A$  in attack expenses. The SA then maximizes its ROI as follows:

$$\max_{T,A} \sum_{i \in T} \left( -C_{atk}(i) + \sum_{j \in A} I(j,i) \cdot T(i) \cdot A(j) \cdot P_s(i) \right) \quad (2.8)$$

Subject to constraints:

$$T(i) \in \{0, 1\} \quad (2.9)$$

$$A(j) \in \{0, 1\} \quad (2.10)$$

$$\sum_{i \in T} (T(i) \cdot C_{atk}(i)) \leq M_A \quad (2.11)$$

Equation 2.8 maximizes profits to a strategic adversary by selecting the set of targets to attack,  $T$ , and actors with whom to share profits,  $A$ . Equation 2.9 and 2.10 constrain the functions as binary variables, and equation 2.11 limits the cost of attack to a particular budget  $M_A$ . These equations can be solved using mixed integer linear programming (MILP). Solving MILP is computationally expensive, and the combination set of actors and targets can become large. However, given that the impact model represents a physical system, some optimizations can be made. The graph can be segmented into regions of operation that contain non-overlapping sets of actors, as commonly found in large interconnected systems, and solved through a divide-and-conquer approach.

The SA may be faced with limited information about the system, as described in Section 2.1.5. When the adversary solves the optimization problem with the perturbed values of the edges and the nodes, it gets a perturbed impact matrix  $I'$ , which is different from the ground truth impact matrix  $I$ . It then bases its attack decision on  $I'$ . In Section 2.2.5, the impact of the imperfect knowledge on the expected gain of the adversary is characterized.

#### 2.2.4 Defense

The defenders are all actors in the system who are fundamentally optimizing their defensive investment decisions. Given the likelihood of an attack  $P_a$ , the likelihood

of an attack being successful  $P_s$ , the expected impact  $I$ , and the cost to defend  $C_d$ , the actor decides to defend a target if  $P_s P_a I > C_d$ . The defensive model is integrated with the other two components, the strategic adversary model and the interdependent impact model, through the parameters  $I$  and  $P_a$ , respectively. The probability of attack is created by the defender's model of the strategic adversary.

## Strategy

Each actor  $a$  in the system owns a subset of targets,  $T_a$ . For each target  $t$ , a binary defense decision  $D(t)$  is made by the owning actor  $a$ .  $D(t) = 1$  means that the asset is defended,  $D(t) = 0$  means it is not. The investment is limited by the defensive resource  $M_D(a)$ . The defender then optimizes as follows:

$$\max_D \sum_{t \in T_a} (P_a(t) \cdot I(a, t) \cdot (1 - D(t)) - C_d(t) \cdot D(t)) \quad (2.12)$$

Subject to the constraint:

$$D(t) \in \{0, 1\} \quad (2.13)$$

$$\sum_{t \in T_a} (D(t) \cdot C_d(t)) \leq M_D(a) \quad (2.14)$$

Equation 2.12 trades the cost of defense against the expected loss due to an attack and results in an optimal defense subject to the constraint in Equation 2.14 which caps the amount of expenditures on defense to  $M_D$ . This can be solved using MILP, as in the strategic adversary case.

## Limiting Information

Similar to the strategic adversary, the defender may have limited information about the system. The impact matrix that the defender bases her decisions on may be formed by a noise-perturbed model of the underlying system, i.e.  $I'$ . The defender is responsible for determining which targets the strategic adversary will attack,  $P_a$ .

This is done by evaluating the SA model from the defender's view of the system. For this, the defender perturbs  $I'$  with her estimate of the knowledge that the adversary has and creates  $I''$ .

### Collaboration in Defensive Strategy

Multiple defenders may wish to coordinate defensive operations for certain targets in the system. Some links may have negligible owner impact but cause substantial losses in other parts of the system. For example, the lowest cost power source becoming disrupted increases costs for all energy buyers, so they may wish to pool resources to defend the low cost source.

Collaboration may occur based on varying levels of agreements. In one extreme, no actors are collaborating, and in another extreme, all actors are collaborating. In order to cooperatively defend a particular asset, all actors interested must have negative impact values for that particular target. At target  $t$ ,  $CD(t)$  is the set of valid cooperating defenders. The optimization is as follows:

Define:

$$C_d(a, t) = \frac{C_d(t) \cdot I(a, t)}{\sum_{i \in CD(t)} I(i, t)} \quad (2.15)$$

Optimize:

$$\max_D \sum_{i \in T} \left( \sum_{j \in CD(i)} (P_a(j, i) \cdot I(j, i) \cdot (1 - D(i))) - C_d(i) \cdot D(i) \right) \quad (2.16)$$

Subject to the constraints:

$$D(i) \in \{0, 1\} \quad (2.17)$$

$$\sum_{i \in T_a} (D(i) \cdot C_d(j, i)) \leq M_D(j) \quad \forall j \in A \quad (2.18)$$

These equations are identical to the earlier set when  $|CD(t)| = 1$ . The optimization in Equation 2.16 makes a decision on the total cost to defend a target when its impact is combined across cooperative defenders.  $P_a(a, t)$  takes into account the fact



that each defender, actor  $a$ , may have a different perceived attack probability based upon the limited information model it uses in assessing defense.

### 2.2.5 Experimentation

#### CPS Model

Refer to Sections 2.1.6 and 2.1.7 for the detailed model explanation.

**Attacker and Ownership Model** The attacker in these scenarios has the ability to reduce a target's capacity to zero. If the attacker targets edge  $(u, v)$  then  $c'(u, v) = 0$  and the impact assessment is done for this perturbation. Although there are several models for attack behavior, this is chosen because it allows for large changes in the system to occur. The impact of attack is measured as independent disruptions to each edge in the graph.

Since this chapter makes no attempt to speculate on the best ownership distributions, a uniformly random assignment of all the assets among the actors is taken. When additional actors are introduced into the model, the number of edges (assets) remains fixed and any one actor's asset pool is subsequently reduced.

#### Experiment 1: Interdependent Model

The focus of this experiment is to analyze the behavior of the interdependent system under attacker perturbations.

The premise of creating a multi-actor impact model is that having multiple actors competing over resources allows for some actors in the system to benefit from attacks. To capture this effect, the summation of positive (and negative) impacts are observed in the system in this experiment. As the number of actors increases, two things will occur. First, competitor elimination becomes more prevalent, i.e., for some functions in the CPS, a monopoly is created, laying the foundation for more profits for some

players. Second, since the attacks are really zero-sum, the gains will be met with corresponding loss potentials.

Figure 2.11 shows the absolute value of gain or loss in the system, averaged across random ownership, versus the number of actors present. The amount of gain in the system increases with actors, as expected, but tapers off as additional competition becomes impossible due to a nearly independent ownership model. The given model has 12 points of competition mapping to the 12 hubs in the gas and electric system, and so saturation occurs around the 12 actor mark in the graph. The takeaway here is that gains are met with losses, and that gains increase with the number of actors.




Fig. 2.11. The total gain and loss in the system, as the sum across impacts felt by all actors, increase as the number of actors in the system increase up to a point of saturation. The sum of the gain and negative loss remain constant.

## Experiment 2: Strategic Adversary

The strategic adversary model is examined to determine what causes most damage to the system.

The SA's goal is to extract profit from the system by attacking assets, subject to a constraint on the total budget she can expend for launching such attacks. The end result of applying costs to attacks is constraining the number of targets or particular targets that the attacker can disrupt. For explorations in this section, the costs are uniform across targets to remove some of the complexities involved in understanding the model behavior and instead a limit to the number of targets will be used.

The SA launches an attack as a set of targets and actors with whom the SA will share in profit, which is determined by solving the optimization function introduced in Section 2.2.3 . To this end, the success metric of the SA is simply the sum of the profits across the target and actor set chosen.

For this experiment, the SA is given a system with varying numbers of actors and varying amounts of knowledge, represented as the standard deviation ( $\sigma$ ) of noise. The intuition is that an increasing number of actors provides a more granular option for target selection. An attack on a particular target may cause, relatively speaking, a gain and a loss to a particular actor. If that actor becomes subdivided into two new owning actors, then the remaining profitable actor can be selected by the SA. The other dimension is that when the SA knows less about the system, through the addition of model noise, suboptimal decisions will be made. Experimentally the SA's target determination is done based on a noisy view of the system, while the actual impact comes from what the ground truth model experiences due to an attack.

Figure 2.12 shows the profitability of the SA, averaged across random ownership distributions, while selecting a maximum of six targets to attack. With a larger number of actors in the system, the success of the SA is increased as expected, with the 2-actor scenario having the worst profitability. This follows the curve in Figure 2.11. As the knowledge level of the attacker is decreased, the effectiveness of the attack also decreases due to poorer decision making.

Figure 2.13 compares the SA's anticipated versus observed profitability. As the knowledge of the SA decreases, and the model becomes noisy, the attacker's anticipated profit does not decrease, but his actual profit does. This suggests a viable




Fig. 2.12. This figure shows the profitability of the strategic adversary versus the amount of knowledge (inverse of noise) that it has about the system. As the noise increases, the profitability decreases. Additionally, as the number of actors increases, the profitability of the SA also increases because of profit opportunities.

defense policy — deception, specifically, making the attacker think that he knows the protected system better than he does in practice. Then, the attacker may be willing to expend greater resources only to realize after launching the attack that he obtained diminished returns (corresponding to the solid line in the figure).

### Experiment 3: The Defenders

The defenders are comprised of every actor in the system, acting in self-interest to mitigate losses due to attacks.

When making assessments about defense, a fixed system budget is assumed (12 assets) and then divided among the actors evenly. This means that in a 12-actor system, each actor can defend a single target, and in a 2-actor system, each actor can defend 6.




Fig. 2.13. This compares the profit of attack for a 6-actor system. The SA anticipated returns, based on the noisy model, do not decay with knowledge level. This means that if the SA is overconfident, the observed returns will be much less than anticipated.

The defender's goal is to minimize the impact of an attack. The metric used for this experiment is the reduction in the impact of the possible attack to the defenders.

To be successful, the defender must accurately reason about the strategic adversary's targets and then move to protect ones which cause a significant loss to itself and are likely to be attacked. This it does under incomplete knowledge (hence the  $\sigma$  for the various parameters that it has to estimate). Further, in estimating the adversary's strategy, it has to speculate on the level of knowledge for the adversary (hence, a speculated  $\sigma$  for the various parameters that the adversary uses). This mechanism is as detailed in Section 2.2.4.

Figure 2.14 shows the effectiveness of defense for a varying number of actors across the noise that the defender has in its model of the system. The Y-axis is the metric that is calculated as follows: compute, for a fixed attack (single asset), the gain to the adversary when the entire system is undefended; compute for the same attack

the gain to the adversary when the defender makes the optimized decision to protect some assets. The metric is the difference of these two values. As the noise increases, the effectiveness of the defense decreases. Intuitively this is because the defender is not completely aware of the impact that an attack has against a particular target and therefore may choose the assets that she wants to defend unwisely. As the number of actors in the system increases, the effectiveness of defense decreases for two reasons. First, the actors are each operating with a smaller defense budget since the funding is constant for the system, thus decreasing per-actor as the actors increase. Therefore, the actor with large negative-impact targets may be underfunded. Second, the actor who should defend an asset may not be the owner, leading to inefficient investing.




Fig. 2.14. The effectiveness of a defense is graded by its impact reduction in ground truth versus the knowledge level of the defender, modeled as noise added to the ground truth. As the number of actors increases, the effectiveness of the defense decreases due to misaligned incentives and a lack of pooled defensive budgets.

Figure 2.15 investigates the impact of collaboration in a system of 4 actors. The collaboration allows the defenders to share in defensive costs, in this case for all assets, as long as they have an aligned defensive incentive. That is, if a target causes damage

to actor A and actor B, A and B will split the defensive costs proportional to their individual impacts. This allows for actors to more optimally defend assets by sharing in costs. This effect wears off as noise increases and the defenders are unsure about which assets are important.




Fig. 2.15. The impact of collaboration is measured by allowing the actors to share in defensive costs. When the costs are shared, more effective investments can be made.

Figure 2.16 compares the impact of collaboration across different actor sizes. In the first case of 2 actors, it is likely that an attack on one target helps actor 1 and hurts actor 2, resulting in a limited collaboration opportunity. In some cases, the attack harms a common supplier or common customer which motivates collaboration. As the number of actors increases, the opportunity for collaboration also increases and results in larger gains. However, for a large number of actors - 12 in our experimental scenario, where there are 96 assets - the incentive for collaboration increases but this is counteracted by forces seen in Figure 2.14 that the effectiveness of defense decreases.




Fig. 2.16. Collaboration allows for actors to improve their defenses. In this case, the system-wide defensive investment is fixed as the number of actors increases, resulting in reduced benefit of collaboration as the number of actors increases and their individual budgets dwindle.

### 2.2.6 Conclusion

In this section, a modeling technique for evaluating cybersecurity defensive investments in interconnected cyber-physical systems was presented. An impact analysis technique enables multiple actors to compete and maximize their individual profits in a flow-optimization problem. The multi-actor approach allows for a strategic adversary to exist who extracts profits from the system by selecting targets to attack and assuming the role of some of the actors in the system. A defensive strategy creates defense optimizations in the face of a strategic adversary. Our experimentation evaluates the impact of attacks, ownership, defensive investments, and collaboration among defenders. It is found that as the number of actors increases and greater competition results, a strategic adversary is able to net more profit from carefully targeted attacks. However, collaboration among actors, even if budget limited, can significantly blunt the effects of such strategic attacks.



### 2.3 Model and Attacks Conclusions, Insights, and Future Work

This chapter presented the first pieces of the dissertation: a model for an example CPS and an approach to attacking and defending it. These pieces, when combined, provide high level insights into the operation of energy markets and how they can be manipulated. The key insight from this work is that economic structures with independent actors can capture strategic adversaries' profit objectives. This model enables new exploration into the energy-based CPS domain by focusing on economic objectives rather than faults and disruptions. Chapter 3 grounds these high level models with realistic negotiation processes. This improves the solution quality by replacing the approximation for multiple actor profits in Section 2.1.5 with an online iterative process.

The solutions presented in this chapter have a few shortfalls that future work can address. First, for simplicity, this model assumes a linear combination of attack profits. In practice, the interdependence of assets may not be linear, and more complex models can better optimize both offensive and strategies. Second, the targets in this work are abstract energy-flow disruptions. Specifically, the attacks constrain energy flows to zero. These targets, however, may have multiple modes of failure that do not necessarily map to a complete stoppage of flow. The model could be improved in future work by analyzing multiple failure modes for each target. In Chapter 4, for example, successful attacks disrupt the ability for energy flows to adjust rather than bringing them down to zero. Finally, this work analyzes a static experiment system. In practice, the system will evolve over time, and the strategy space for attack and defense may not exist at a single optimal point. Future work could introduce a temporal aspect to the strategy space. Tangentially, non-energy CPS may have different objective functions, but many of the techniques presented in this chapter can still be applied to different domains simply by replacing Equation 2.1 with another objective. Future work could also address different domains.

The next chapter introduces more technical models for energy system operation. In these new models, communications exist between market players. Since these communications occur in real-time, latency and temporary network outages can influence market behavior.

### 3. ENERGY CPS OPERATION

This chapter covers two technical market systems that use cyber-physical systems to negotiate the usage of energy. The techniques presented here allow the targets in Chapter 2 to be mapped to low-level network attacks. With the work presented here, an adversary can attack a realistic smart grid deployment and extract additional profits from the system, as later demonstrated in Chapter 4.

Section 3.1 is based on joint work with Dylan Shiltz and Thomas R. Nudell from MIT’s Active Adaptive Control Laboratory. Sections 3.1.2, 3.1.4, and 3.1.5 and the dynamic market mechanism (DMM) model are primarily the contribution of those authors. The communication modeling in Section 3.1.3, experimental setup, and experimental results and analysis are novel contributions for this dissertation. The second half of this chapter (Section 3.2) is a market method created for analysing real-time CPS.

#### 3.1 A Framework for Evaluating the Resilience of Dynamic Real-Time Market Mechanisms

##### 3.1.1 Introduction

High penetration of renewable generation introduces challenges in many areas of power system operations, including real-time market operations. In particular, the fluctuations in renewable generation motivate a need for real-time coordination of distributed energy resources, which can respond rapidly to intermittent generation. A promising approach to facilitate this coordination is through dynamic economic dispatch algorithms such as those proposed in [3, 36–43], paired with flexible consumption through DR and flexible storage resources. These mechanisms enable ef-

efficient integration of unpredictable generation and flexible loads into the power grid with the goal of improving grid resilience and operational efficiency. Naturally, real-time coordination of distributed energy resources, especially over a large geographic region, requires communication across networks [44]. Such networks can experience non-negligible latency, loss of information, and congestion. Dynamic market mechanisms and flexible consumption and storage will ideally improve grid performance; however, less is known about their resilience behavior in the face of delays, outages, and other disruptions to the underlying communication infrastructures, which may be unavoidable. In this section, a framework is proposed for evaluating the resilience of dynamic real-time market mechanisms for managing electric power grids in the face of realistic network and power disruptions. This framework may be used to evaluate the resilience of any number of market mechanisms (such as those described in [3, 36–43]) operating on various communication infrastructures and power system networks.

In typical planning and operation of the power grid today, several market-based control layers are already employed to maintain a balance between generation of and demand for electricity. At the slowest time-scale, generation units are committed well in advance, bidding to supply power in day-ahead markets (DAM). This creates a tentative dispatch schedule for generators. Real-time markets (RTM), which operate on the time-scale of minutes, revise this dispatch schedule to accommodate changes in supply, account for inaccuracies in the predictions of load, and adjust to variable renewable energy resources (RER). These markets also serve to reduce outage-induced stress and inefficiency during power shortages or surpluses. Much of the intermittency associated with renewables typically occurs at this faster time scale, necessitating dynamic algorithms, including dynamic real-time market mechanism and price-based coordination of DR.

In the current market structure, most loads are not price-sensitive (i.e. they are price-setters rather than price-takers). For example, many residential consumers pay a fixed rate for power that is set by utility companies and regulated by local and/or federal agencies; these rates change, at the fastest, over the course of months. The

DAM and RTM prices can potentially exceed these fixed rates during peak demand or supply shortages, disrupting market efficiency. Traditionally, the market cannot operate in a way that provides maximum *social welfare* to all participants because of these traditional inflexible loads. In contrast, dynamic market mechanisms allow flexible consumers to act as price-setters rather than price-takers, thereby allowing near real-time negotiations of prices, generation, and flexible consumption set-points which in turn can lead to optimized social welfare. Implementing a dynamic real-time market mechanism, however, entails substantial communication overhead for exchanging real-time information and iteratively negotiating prices.

In most existing algorithm designs [3, 36–43], the communication layer is typically idealized or ignored during market method evaluation. In reality, however, communication networks cannot always operate with strict guarantees on latency or disruption. Routing issues, link outages, and message loss are commonplace in large networks due to the number of devices and their configurations. Large scale outages on commercial networks are not uncommon [45]. Models of dynamic markets that do not consider such network issues may prove unreliable or inconsistent when implemented in a realistic communication environment that is prone to such outages. Therefore, the influence of the dynamics of communication networks on markets and their convergence, when driven by faults and failures, needs to be analyzed in detail before these methods can be adopted widely on the smart grid. The framework proposed in this section can be used for this purpose—to evaluate the resilience of these market mechanisms and better understand how they may handle contingency events in the power system under constraints of realistic communication implementation.

One strategy to account for realistic communication constraints is to design *latency aware* algorithms, which assume a non-zero latency required for communication. For example, DYMONDS [46] requires that communications be designed such that all of the information required for operation is available in the central control location. It is perfectly rational to design a latency aware algorithm in principle, but it may be difficult to guarantee these specified latencies in practice. Hence, any latency aware

algorithm can also benefit from the proposed framework by evaluating its viability when the communication network does not operate as designed.

Several prior works evaluate the susceptibility of broader cyber-physical systems to cyber-attacks. Work in [47], for example, focuses on cybersecurity aspects in confidentiality and integrity, but it does not focus on the impact of network latency on system operation. Work in [48] evaluates energy systems with distributed resources—the target system for work in this section. Their security focus, however, is on integrity attacks via spoofed messages and insecure communications. In systems with distributed energy resources, however, the communication network interconnections may span large distances, and thus suffer from data availability concerns that have not been addressed. The framework provided in this section focuses on the impact of information availability on power grids, with an emphasis on market operation.

### **Framework Elements**

The proposed framework consists of the market layer, including the decision making components (ex. an independent system operator (ISO)), a communication layer (ex. clients and servers), and the physical layer (ex. generators, loads, and power lines) as shown in Fig. 3.1. The communication layer maps the agents in the physical network—generators and flexible loads—to nodes in the communication network. The market layer represents decisions of economic dispatch where a balancing authority (ex. ISO) and market participants (ex. generators and flexible consumers) communicate with each other. In practice, the particular architecture of any of these three layers may vary. This layered structure enforces causality across communication boundaries and adds realism to the market mechanism design.

Using this framework, a standard scenario is considered by introducing faults at the physical layer and failures in the communication network. Electrical power systems must be designed and operated such that acceptable performance is maintained following a contingency. The North American Electric Reliability Corporation

(NERC) defines a *credible* contingency as one that is both plausible and likely. The failure of any single element of the power system may be considered a credible contingency (sometimes referred to as  $N - 1$  contingency), while the simultaneous failure of multiple elements that are not physically or electrically related is not likely, and is therefore not a credible contingency [49]. In this section, the most common type of single element failure is considered: generator outages.

Existing control loops are designed to stabilize the power system following such disturbances. However, today’s traditional real-time markets (RTM), which determine dispatches on the order of 5 minutes, often clear some time before the operating hour, and at best, can take several minutes to re-allocate generation. In contrast, dynamic real-time market mechanisms can enable the grid to adapt more quickly, responding to on-line information, making the grid more resilient to disturbances. These new methods themselves, however, rely more heavily on communication infrastructure. The resilience study developed in this section attempts to evaluate the ability of such mechanisms to respond to common contingencies within an adverse communication environment.

A resilient mechanism will adapt to changes or threats to the system and still be able to match supply and demand as closely as possible at all times. For this purpose, metric is introduced that captures the impact of a physical disturbance and the ability of the market mechanism to adapt to this physical contingency and reallocate power in the face of simultaneous disruptions to the communication network. Hence, this metric can be used to evaluate the resilience of any particular market mechanism with a given communication infrastructure and electric power system model and contingency.

The remainder of the section is organized as follows. Section 3.1.2 introduces dynamic electricity market mechanisms. Section 3.1.3 models the operation of these market mechanisms over realistic communication networks. Section 3.1.4 formally defines the resilience metric. Section 3.1.5 details the experimental scenarios, which start with a 118-bus test case, used to illustrate the framework. Section 3.1.6 discusses

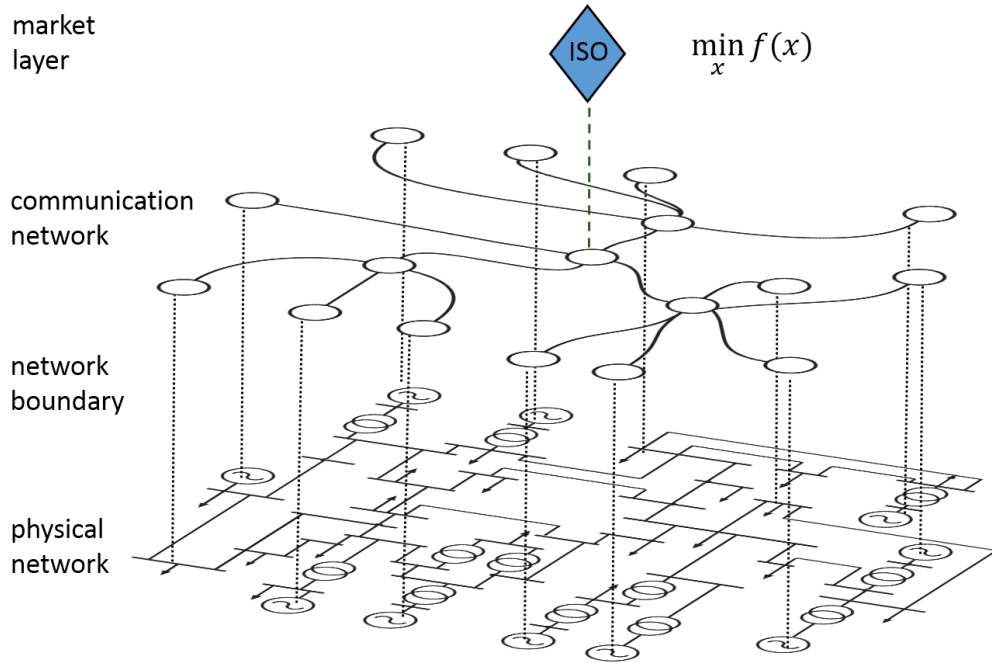


Fig. 3.1. Each agent in the physical network—generators and flexible loads—is associated with a client in the communication network, these clients send and receive information to and from the server at the independent system operator (ISO) who operates the market.

simulation results, and demonstrates the utility of the framework for evaluating the resilience of dynamic real-time market mechanisms. Section 3.1.7 provides concluding remarks.

### 3.1.2 Dynamic Real-Time Market Mechanisms

An electric power grid is a network of high voltage AC transmission lines that route power between generators and consumers over large distances. Today this network is managed through a hierarchy of electricity markets, in which generators and consumers can negotiate the price and quantity of power at various locations in the grid. In deregulated electricity markets, an economic dispatch is performed on



the order of every 5 minutes. This dispatch, calculated in a centralized fashion by an independent system operator (ISO), determines the most cost effective way to meet power demands subject to constraints on generators and the grid.

The introduction of intermittency and uncertainty into the electric power grid through the widespread adoption of RERs, however, has caused the traditional markets to become decompensated. Managing RERs, which commonly suffer from intermittency on a very fast time-scale, and the potential of adjustable demand through DR resources necessitate a dynamic framework. The former introduces issues of strong intermittency and uncertainty, and the latter a feedback structure where demand can be modulated over a range of time-scales. Both of these components dictate a new look at market mechanisms with a control-theoretic perspective. In smart grid literature this has sometimes been called *price-based control* (see [50] and references therein) or *transactive control* (see [51] and references therein). Throughout this section the underlying algorithms of these methods will be referred to as *dynamic real-time market mechanisms*.

Dynamic real-time markets represent expanded participation opportunities in the market process, which has been enabled by technical innovations that allow for rapid communication among a wide, automated user base. This allows a market to efficiently embrace RERs and DR with real-time decision making but requires large digital communication networks. In turn, frequent communication enables grid components (clients) to quickly adapt to changing conditions. It also allows for more accurate predictions of renewable generation and load fluctuations, as decisions are made closer to real-time.

## Market Mechanism Operation

Although electricity market mechanisms vary in their implementation, many of the dynamic real-time mechanisms share a few common features. Each generator typically has a *cost curve* that describes the marginal cost of generation as a function

of the power being generated. These curves tend to be convex (i.e., as a generator approaches its maximum operating limit, its marginal cost of generation increases). Similarly, each flexible consumer has a *utility curve* that describes the marginal utility of consumption as a function of the power consumed. It is assumed that these utility curves are concave [3, 36–43]. In addition, all generators and consumers have limits on how much power they can produce or consume at any given time. A concrete implementation with quadratic cost and utility curves is shown in Section 3.1.5.

The objective of a market mechanism is to maximize system utility and minimize system cost, subject to constraints of the market players as well as the physical transmission system. This can be written as an Optimal Power Flow (OPF) problem, which is an optimization problem of the form

$$\min f(x) \tag{3.1}$$

subject to

$$h(x) = 0 \tag{3.2}$$

$$g(x) \leq 0 \tag{3.3}$$

where  $x$  is a vector of system states (including generation, consumption, voltage angles, etc.),  $f$  is a cost function of these states,  $h(x)$  is a set of equality constraints enforcing power balance at each bus, and  $g(x)$  is a set of inequality constraints enforcing bounds on generation, consumption, and line transmission. The most general form of (3.1) is a fully non-linear AC OPF that is non-convex and NP-hard. Thus, for real-time operation it is common to use a linearized DC OPF formulation [52], which

neglects reactive power flow and AC line losses. This allows the equality constraints  $h(x)$  to be written as a linear function

$$h_n(x) = \Delta_n + \sum_{i \in \mathcal{D}_n} P_i - \sum_{i \in \mathcal{G}_n} P_i + \sum_{m \in \Omega_n} B_{nm}(\delta_n - \delta_m) = 0 \quad \forall n \in \mathcal{N} \quad (3.4)$$

where  $\mathcal{D}_n \subset \mathcal{D}$  and  $\mathcal{G}_n \subset \mathcal{G}$  are sets of flexible consumers and generators at node  $n$ ,  $\Omega_n$  is the set of nodes adjacent to node  $n$ ,  $\mathcal{N}$  is the set of all buses in the network, and  $B_{nm}$  are the susceptances of the lines from node  $n$  to node  $m$ . Each bus also experiences a conventional (inflexible) demand  $\Delta_n$ .

If the objective function  $f(x)$  is convex, then Problem (3.1) is convex and can be solved efficiently. Price-based solutions of (3.1) usually involve calculating locational marginal prices (LMP's), denoted by  $\lambda$  at each bus in the system. The power produced or consumed by each market player is denoted by  $P_i \in x, i \in \mathcal{V} = \mathcal{D} \cup \mathcal{G}$ , and is typically some function of  $\lambda$  and the grid state  $x$ . That is,

$$P^{k+1} = \psi(x^k, \lambda^k) \quad (3.5)$$

where  $\psi$  may contain information regarding grid topology, cost and utility curves, grid frequency, and other network parameters. In (3.5), the superscript  $k$  denotes the iteration index, with  $z^k$  denoting the value that  $z$  takes at time  $t_k$ , and  $P$  denotes a vector with its  $i$ th element given by  $P_i$ . The power setpoints are iteratively updated with period  $T_k$ , such that  $x$  converges to  $x^*$ , the optimal dispatch for the grid. It should be noted that price-based mechanisms are not the only proposals in the literature (other approaches include dynamic programming, integer programming, and other non-derivative methods), but in this section is focused on price-based methods.

The structure of (3.5) suggests a decentralized execution over a large digital communication network with synchronized updates. If information is successfully sent and received during the negotiation window  $T_k$ , then the algorithm works as de-

signed. However, if communication is lost or is delayed, the algorithm suffers from network impacts which may adversely affect its performance. In order to understand the network impacts, the next part of the framework describes communication models.

### 3.1.3 Modeling Latency in Communication Networks

Solving the optimization problem (3.1) using a dynamic real-time market mechanism requires communication between widely distributed parties—namely clients (ex. loads and generators) that may be several hundred miles apart and a server (ex. ISO). To achieve a near real-time synchronization between the loads and generators, the market model must be fast enough to converge with reasonable computation power, and the interaction between the clients and server must occur via a realistic communication network.

Most market mechanisms operate by iteratively updating price in real time in response to changing physical conditions such as increased load. These physical conditions occur at the ends of the network and must be communicated with the central market coordinator (ex. ISO) to establish optimal prices. Once new prices are available, these too must be communicated to the clients. When network conditions are non-ideal, this communication process can become irregular, disrupting normal market behavior.

### Iterating Markets Over Networks

Market mechanisms are designed to find solutions to the optimization problem (3.1), often using iterative solution techniques such as the Newton-Raphson [3] or interior point [53] methods. Network-agnostic implementations of these methods assume that the computations and iterations occur instantaneously and synchronously in the system. On a practical computer network, however, each iteration takes some non-zero time to complete. The numerical solution method is broken into discrete




Fig. 3.2. Relevant time scales for distributed market mechanisms

steps with a real-time duration of  $T_k$ , and the maximum number of allowed iterations to converge is  $\frac{T_s}{T_k}$ , with  $T_k \ll T_s$  (see Fig. 3.2). The market mechanism sends a negotiation  $x^k$  and waits  $T_k$  seconds for a reply from the clients. If every client replies within  $T_k$  seconds, then the algorithm operates as designed. If no reply is received, however, the algorithm may enter an undefined operational state and can suffer network impacts. The rest of this section defines how the market will operate in the presence of irregular communications.

### Latency Model

Latency is introduced to market mechanism algorithms by delaying the round-trip interaction between the central market and the distributed clients, as a perturbation of an iterative optimization algorithm. In this way, the control actions taken because the market outcome are delayed [54, 55]. Each client  $i$  has a pair of states  $x_i, \lambda_i$  that is delayed by a latency function  $\tau_i(k)$ . This shift must be negative to maintain

causality, and assuming time starts from zero,  $k - \tau_i(k) \geq 0$ . Delay is introduced to (3.5) as

$$P^{k+1} = \psi(x^{k-\tau(k)}, \lambda^{k-\tau(k)}) \quad (3.6)$$

where  $x$  and  $\lambda$  are defined in (3.5) and  $\tau(k)$  can be thought of as a vector of the latencies of each client.

An example of a latency model applied to client  $i \in \mathcal{V}$  can be expressed as

$$\tau_{c,i}(k) = \begin{cases} k & k < \ell_i \\ \ell_i & k \geq \ell_i \end{cases} \quad (3.7)$$

where  $k$  is the time-index and  $\ell_i$  is a constant-latency value for client  $i$ . The simplified constant latency model in (3.7) is applicable to network conditions where each client has a fixed minimum latency to communicate with the ISO server e.g. speed of light limitations due to grid geography.

Another model shown in (3.8) captures network outage situations, and we denote this function as  $\tau_x$ . In this case, a set of clients, denoted  $\mathcal{X} \subset \mathcal{V}$ , is unable to communicate with the market. Formally,

$$\tau_{x,i}(k) = \begin{cases} k & i \in \mathcal{X} \\ 0 & i \in \mathcal{V} \setminus \mathcal{X} \end{cases} \quad (3.8)$$

The client is unable to receive an updated price, and the market evolves as if the client is unable to change its power level. The model (3.8) assumes the outage occurs at  $k = 0$ , although the value of  $\tau_{x,i}(k)$  can be shifted to the outage start time. The model (3.7) and (3.8) are a compact representation of the latency models used in [54, 55].

For simplicity it is assumed that the market mechanism uses a zero-order hold during non-communication periods in both (3.7) and (3.8). That is the part of the state  $x$  controlled by client  $i$  does not advance. It is also noted that (3.7) and (3.8)

are primitive models which can be combined to create more complex scenarios that include both disconnects and constant latency.

### Practical Networking Constraints on Market Mechanisms

Up to this point,  $\tau$  has been modeled as an integer parameter, offsetting  $k$ . Real computer networks, however, experience real-valued latency  $L \in \mathbb{R}_{\geq 0}$ . The mapping from real latency  $L$  to  $\tau$  can be represented as

$$\tau_i(k) = \left\lfloor \frac{L(t_k)}{T_k} \right\rfloor \quad (3.9)$$

where  $\tau_i(k)$  is the delay shift at time step  $k$ ,  $t_k$  is the mapping between the step  $k$  and time, i.e. as  $t_k = kT_k$ ,  $L(t)$  is the real-valued latency at time  $t$ , and  $T_k$  is the step size.

The convergence rate of stochastic gradient descent based methods is bounded by  $O(\sqrt{\frac{\tau}{T_k}})$  [56, 57]. Comparing  $L(t) = 2T_k$  to  $L(t) = T_k$ , for example, gives some intuition to why this is true. In a network-agnostic implementation,  $\tau$  may be forced to 0 by increasing  $T_k$  to absorb the latency difference. As a result, the convergence rate would be half. The  $\tau$  model, however, operates like a pipeline where new iterations may be based upon stale values. As a result, decreasing  $T_k$  with a fixed  $L$  provides diminishing improvements in solution quality.

The communication systems of the future smart grid have not yet been finalized [58], and there is a direct relationship between reliability and cost. Internet-scale communication, for example, relies on a best-effort model to keep costs low. This means that there are no guarantees on delivery, bandwidth, or latency among clients connected to the Internet. While many competing models in the past such as circuit switching networks provided dedicated, rigid bandwidth and latency guarantees, they were ultimately too expensive compared to best-effort methods since they eliminated multiplexing opportunities (the use of idle bandwidth). Establishing high-reliability latency-guaranteed networks between power generation stations, often located in re-

mote areas, presents a high-cost barrier to reliable communication assumptions. For this reason, best-effort networks may provide a cost-optimal solution, as long as the market mechanisms can sustain operation in poor networking environments. If design shows that there are significant benefits to reliable latency values, then more rigid network constraints could be implemented such as those for the IEC 61850 technical standard that requires performance minimums for certification.

In order to understand the consequences of the latency models in the above discussion in the context of resilience, the next section develops the final component of the framework: a resilience metric.

### 3.1.4 Resilience Metric

The primary responsibility of the dynamic real-time market is to allocate resources to match supply and demand repeatedly over very short time-scales. The dispatched power matches demanded power as closely as possible at all times, but primary control systems and secondary markets will always compensate for any discrepancies. In other words, the amount of power that cannot be allocated by the market mechanism must be accounted for by other means, either through costly ancillary services or through adverse impacts to the grid such as frequency deviations from 60 Hz. The effectiveness of the mechanism can therefore be measured based on the mismatch between dispatched power and actual power generated at a particular time throughout the system, termed as residual power, denoted  $R_p(t) \in \mathbb{R}$ . The resilience metric that is proposed in this section, therefore, is a measure of  $R_p(t)$ .

In the current context, it is argued that a resilient mechanism should adapt to changes or threats to the system and still be able to match supply and demand as closely as possible at all times. A perfectly resilient mechanism will have an  $R_p(t)$  that is always zero, that is zero impact, even in the face of disturbances. A non-resilient



system can then be viewed as a system with non-zero  $R_p(t)$ , with its magnitude indicating the degree of non-resilience. Therefore, the resilience metric is defined as

$$R(t) = 1 - \frac{|R_p(t)|}{R_p^{max}} \quad (3.10)$$

which is a non-dimensional quantity with  $0 \leq R(t) \leq 1$ . Notice that  $R(t) = 1 \forall t$  denotes a perfectly resilient system and  $R(t) = 0 \forall t$  indicates a perfectly non-resilient system. The metric in (3.10) captures the impact of failure, which is one key component of risk modeling of cyber-physical systems [23, 48, 59].

Other aspects of risk modeling include threat and vulnerability, which capture the likelihood of a particular failure in both the physical and communication networks [23]. Vulnerabilities include failures of the advanced metering infrastructure and manipulation or disruption of SCADA systems [60]. Developing a more comprehensive resilience metric, in comparison to (3.10), which additionally captures threat and vulnerability is still an active area of research and outside the scope of this section.

Using (3.10) along with realistic communication, computation, and market mechanism models allows us to evaluate and compare the resilience of future grid operating strategies and provides insight into overall system performance. A concrete example of such future grid operation consisting of the Dynamic Market Mechanism (DMM) developed in [3] used to manage the IEEE 118-bus test system is described in the proceeding sections. The experimental setup is presented next, followed by simulation results in Section 3.1.6.

### 3.1.5 Experimental Setup

In this section, the experimental setup that is used to evaluate the resilience framework is explained. Three experiments are designed that evaluate the resilience of the DMM described in (Section 3.1.2) under different communication network stresses. Each experiment uses the same physical system, the IEEE 118-bus. At the beginning of each experiment the highest-output generator is removed from the system. The

resilience of the DMM is measured by  $R(t)$  defined in (3.10) with  $R_p^{max} = \bar{P}_i$  of the removed generator. All simulations were carried out in MATLAB<sup>1</sup> using the Matpower 118-bus test case [61].

### Physical System Scenario

The experiment begins with the IEEE 118-bus test system [61] with a suite of loads—two types of dispatchable DR along with conventional loads—and generators connected to various buses throughout the system. The conventional generators and loads are defined by the 118-bus test case [61]. There are fifty-four dispatchable generators in this model. Additionally, a total of thirty dispatchable demand-response units are distributed arbitrarily throughout the system with power consumption limits of  $[10, 60]$  MW. The thermal limits on the transmission lines are assumed to be 300 MW such that the system is partially congested.

Cost and utility curves, mentioned in Section 3.1.2, used in this example are quadratic curves with an additional barrier function [62] to accommodate minimum and maximum bounds on power generation or consumption. Specifically, the cost and utility curves are expressed as

$$C(P_i) = b_i P_i + c_i P_i^2 + \frac{M}{(P_i - \bar{P}_i)^2} + \frac{M}{(P_i - \underline{P}_i)^2}, i \in \mathcal{G} \quad (3.11)$$

$$U(P_i) = b_i P_i + c_i P_i^2 - \frac{M}{(P_i - \bar{P}_i)^2} - \frac{M}{(P_i - \underline{P}_i)^2}, i \in \mathcal{D} \quad (3.12)$$

where  $b_i$  and  $c_i$  are base and incremental cost/utility parameters for the  $i$ th agent, respectively, and  $\bar{P}_i$  and  $\underline{P}_i$  are maximum and minimum generation or consumption. The parameter  $M$  determines the steepness of the barriers, and in the experiments  $M = 5$ . The cost parameters for the generators are defined in [61]. DR cost parameters are selected from the uniform distribution  $b_i \in \mathbf{U}(0, 45)$ ,  $c_i \in \mathbf{U}(-2, -1)$  for  $i \in \mathcal{D}$ .

---

<sup>1</sup>MATLAB is a registered trademark of The Mathworks, Inc

## Summary of DMM

The dynamic market mechanism used here was originally proposed in [3], and it uses a Newton-Raphson-like primal-dual interior point method [62] to update (3.5). The decision variables are  $P_i, i \in \mathcal{V} = \mathcal{G} \cup \mathcal{D}$ , and  $\delta$ , which denote generation, flexible consumption, and voltage angles respectively. We denote the state vector as  $x = [\delta^T P^T]^T$  and the objective function as  $f(x)$ . The Lagrangian of the optimization problem is  $\mathcal{L}(x, \lambda) = f(x) + \lambda^T h(x)$  where  $h(x)$  denotes power balance at each bus in terms of the state variables. The DMM iterates can be expressed as

$$x^{k+1} = x^k - \alpha(H)^{-1} \nabla_x \mathcal{L}(x^k, \hat{\lambda}^k) \quad (3.13)$$

$$\lambda^{k+1} = \hat{\lambda}^k - \alpha h(x^k) \quad (3.14)$$

where  $\hat{\lambda}^k = M_1(h(x^k) - M_2 \nabla f(x^k))$ ,  $H$ ,  $M_1$ , and  $M_2$  are constant matrices, and  $\alpha$  is a positive step size. See [3] for specific details. In our experiments  $\alpha = 0.01$ .

## Networking Scenario

The communication network is a hub and spoke model where each client, either a generator or a DR unit, has a direct link to the ISO who facilitates the market. The nominal network scenario can be described by a particular  $\tau$  function, as mapped through (3.9). For these experiments, a DMM with negotiation iteration length of  $T_k = 30$  ms is used as specified in [3]. Therefore  $\tau = 1$  maps to  $L \in [30, 60]$ . Link outages are modeled as in (3.8), with the size  $|\mathcal{X}|$  indicating the number of links disrupted. In practice, multiple links could be targeted, representing a strategic attack on the network. However, analysis of such attacks is outside the scope of this section. Three networking scenarios are described next.

**Experiment 1: Single Link Outage** For a single link outage  $|\mathcal{X}| = 1$ . This simulates a single fiber cut, for example, or other sudden network failures. In total, 84 outage simulations are evaluated, one for each client.

**Experiment 2: Small Constant Latency** Small constant latencies represent minimum communication delays for deployed systems. These delays arise due to the speed of light, the actual lengths of communication cables, the store-and-forward delay of intermediate networking equipment such as routers and switches, and the technology in use (fiber, radio transmission, etc.). Part of the system design is selecting a communication rate for processes that use iterative solution techniques. If this communication interval is too short, then the response of the system will become degraded. Design values for latency may be exceeded with geographic growth of the control area, backup network links, or general system design.

For Experiment 2, the latency of every client is identical and scaled from  $L = 0$  ms to  $L = 120$  ms in 30 ms increments using model (3.7) with the parameter  $\ell_i$  ranging from 0 to 4 (i.e.,  $\frac{L=0ms}{T_k=30ms} = 0, \dots, \frac{L=120ms}{T_k=30ms} = 4$ ). In contrast to Experiment 1, no links are taken offline in Experiment 2.

**Experiment 3: Non-Uniform Latency** Many cases can arise in networks where there is a non-uniform latency across the different client to server connections. For example, some clients are farther geographically from the server than others and are therefore more susceptible to latency than those that are closer. Additionally, some clients may use older communication technologies—especially in backup situations—that have high latencies such as satellite-based communications.

To capture this, the model (3.7) is used, as in Experiment 2, but with non-uniform  $\ell_i$  across clients. This non-uniformity is introduced via independent Gaussian noise. Each client’s latency is assigned from a truncated Gaussian with values between 0 and  $2\mu$  as  $\mathbf{N}(\mu = 60 \text{ ms}, \sigma)$ , where  $\sigma$  is the controlled parameter. This creates a situation where some clients are able to respond more rapidly than others during the simulation.




Fig. 3.3. The system responds to a worst-case single-generator failure. The market minimizes residual power in two identical experiments, but one has a communication link outage to a single client—the one that causes the most impact.

### 3.1.6 Experimental Results

#### Experiment 1: Residual Power with Single Link Outages

In this experiment, every network link between the ISO and the client (generator or demand-response agent) is independently disrupted for a total of 84 simulations. The cumulative residual power for each outage is used to rank the links in order of impact. Fig. 3.3 shows the response to the power disruption for both the baseline (outage-free) and worst-case link outage scenarios.

The results indicate that this particular market mechanism is resilient to single link outages. In the worst case, a single link outage is identical to removing two market players from the system for this scenario (one by power fault and one by communication fault). As long as the underlying power system has the capacity to adjust to these combined worse case situations, then isolated link outages will not adversely impact grid resilience.

In fact, one could argue that this dynamic real-time market mechanism has made the power system more resilient at this fast time scale, despite the isolated link outages. To understand this, consider the RTM used today, which will clear for 5 minute intervals over an operating hour at least 60 minutes before the start of that operating hour. This means that the market would make no adjustments over the course of the simulation. Instead, the system operator must rely on ancillary services such as AGC, or rely on corrective action—essentially a manual override implemented by the central ISO. In contrast, the DMM has adapted to the physical fault and reallocated most of the power optimally.

While the framework demonstrates that the DMM response is satisfactory, the algorithm is unaware of its networking environment. At convergence, a small amount of power has not been reallocated by the market mechanism. This amount is proportional to the remaining difference between dispatched and actual power at the client suffering a link outage. This quantity can be minimized by tuning DMM parameters, namely  $\alpha$  and  $M$ . Alternatively, if the algorithm were aware of the network failure, then it could drop the affected client from the market and dispatch the power to other market participants. This indicates practical design challenges in market mechanisms of detecting and perfectly adapting to network disruptions.

## Experiment 2: Small Constant Latency

Fig. 3.4 shows the impact of low constant latency values on algorithm performance. As latency increases, the response of the system becomes slower. In an integrative process, the effective gain of the control system is proportional to the rate of communicated updates since the state uses a zero-order hold in between update periods. When the latency surpasses the design threshold of  $T_k = 30$  ms, the impact is similar to a reduction in the  $\alpha$  gain parameter in update equations (3.13)–(3.14). While the recovery of the test system slows down as the latency increases, the resilience is still acceptable. In particular, DMM is still able to fully adapt to the loss of a




Fig. 3.4. The mean latency for each client is increased. The higher latency values reduce the rate of communication, and the convergence speed of the system for different latencies is shown.




Fig. 3.5. The variance in latency is increased for a small mean value. After the system apparently stabilizes, the response begins to oscillate.

generating unit—even the highest constant latency case that was tested—well before a traditional RTM would be able to take any action.

### Experiment 3: Non-Uniform Latency

Non-uniform latency can create instability in the DMM, as Fig. 3.5 illustrates. Initially, the system responds as expected with the resilience metric driving towards unity. When the system nears convergence (after 30 seconds), the resilience begins to deteriorate for all distributions of non-uniform latency. The oscillations in the  $R(t)$  signal indicate instability in residual power due to commanded power levels. This is because the utility and cost curves are non-linear around their upper and lower bounds. Small changes near the bounds result in large jumps in the cost gradient values that the DMM algorithm uses to optimize social welfare. Since some clients can respond faster than others, many iterations can pass with these inflated gradient values. As a result, the system moves much farther away from the power constraints than it otherwise would. This behavior ultimately drives the system to instability after sufficient time and crashes the market.



An interesting note is that the distributions simulated here have a fairly small variance from the mean. If larger variances are used, then the instability could occur even more quickly. These results indicate that asynchronous iteration across market players can lead to unstable operation, and that market designers should be aware of the impacts of non-uniform latency.

In summary, the framework shows that the DMM is tolerant to single link outages and small uniform delays. In these scenarios, nearly all of the power lost by a worst-case generator trip is reallocated within five minutes. Hence, the DMM can make the 118-bus test system resilient to  $N - 1$  contingencies despite non-ideal communication network conditions. The framework has also revealed that non-uniform latencies, even those with a small mean and variance, can lead to instability in the market mechanism. These procedures and tests could be repeated on a different power system model with a different market mechanism and/or communication network.

### 3.1.7 Conclusions

In order to efficiently embrace DR resources, the smart grid is steadily moving toward dynamic, real-time markets that require frequent and reliable communication over wide area networks. Yet, little is known about how such markets may perform in real-world implementations as the underlying communication networks may be subject to latency and faults that can impact the market performance and ultimately the performance of the physical power grid. In this section, a framework was developed for analyzing and evaluating the resilience of real-time markets in the face of latency and faults in the communication network. The framework may be used as a starting point to evaluate the resilience of any power system operating with any given market mechanisms implemented with a realistic communication network. The resilience metric developed captures the impact of physical failures and the ability of the market mechanism to adapt to and recover from these failures in an adverse communication environment.

To illustrate the framework’s utility, a recently proposed dynamic market mechanism was simulated with an IEEE 118-bus test system that was subject to a  $N - 1$  contingency scenario of the physical grid. The framework illustrated that the system was resilient to single communication link failures as well as constant low-latency scenarios. This is a significant resilience improvement enabled through the fact that the DMM was able to respond at a much faster rate than the typical time scales of current RTMs. The framework also provided insight into some of the potential vulnerabilities of the DMM. In particular, when there are non-uniform latencies, the DMM may result in oscillatory behavior leading to instability. This in turn suggests that attention should be paid to designing a resilient DMM that accommodates such latencies.

## **3.2 Dynamic Pricing for Smart Grids in the Presence of Non-Linear Network Conditions**

### **3.2.1 Introduction**

Modern electric grids and distribution systems are plagued by pervasive inelastic loads that do not respond to changing grid conditions such as the unpredictable availability of renewable energy resources (RER). Therefore, when market conditions indicate scarcity, these loads happily consume energy, agnostic to soaring energy-source costs. Unregulated power markets suffer from large, dramatic swings in prices due to this inelasticity especially when unpredictable or transient power conditions arise. These swings, absorbed by power producers, intermediate brokers, and distribution system operators, cause serious financial losses, drive companies such as Duke Energy out of competitive markets [63], and result in grossly inefficient uses of electricity. A study for the NY-ISO, the independent (power) system operator (ISO) for the state of New York, indicates benefits from elasticity could be as high as \$403 million per year [64]. A solution that introduces elasticity exists—demand side man-

agement (DSM) in the smart grid (SG)—but it presents a unique set of networking and optimization challenges that have not yet been solved.

Creating flexibility requires both exposing consumers to an incentive signal and enabling them to act on it. Demand side management [65] embodies the set of techniques and technologies that enable load elasticity. This section is specifically focused on the transactive control elements—those which enable elasticity. For example, a technique may negotiate a price for power, send it to a consumer, and allow the consumer to adjust a load via intelligent devices such as an Internet-enabled thermostat. Most existing power markets operate with large bulk producers that are unable to respond to rapid changes in price signals. For example, the NYISO estimates that 96% of energy is cleared in their day-ahead market with only 4% of energy transactions occurring in their 5-minute real-time market. At the same time, however, distribution level system operators (DLSO) are burdened by rooftop solar panel integration [2], charging penalties due to the disparity in net metering revenues and actual market costs. DSM techniques are needed that can operate at the distribution level with large numbers of individual consumers. Existing techniques, however, rely on predictable, cooperative consumers and real-time information exchange to operate successfully. In this section a technique for distribution-level transactive control is presented and evaluated under non-linear network conditions.

DSM techniques can be broadly split into two categories: transactive control and direct load control (DLC). In DLC, a central authority creates elasticity by directly modifying the energy set points for various loads in the grid with the permission of the user under a contract. An aluminum smelting plant may be asked to shut down for the day or a clothes dryer in a home may be temporarily disabled, for example. Techniques such as demand response (DR) and agile balancing [43] fall in this category. These techniques, however, centralize control and eliminate consumer’s choices—choices that often change in real-time. To account for time variations, a transactive layer is added to DR via dynamic pricing signals. Instead of directly controlling loads, the central

authority sends dynamic price signals, reflective of market conditions, and consumers respond to these prices.

Several techniques and methods exist for managing price signals, but they have two shortcomings that are addressed by work in this section. First, most techniques, such as those in [41, 53, 66], rely on optimization techniques such as the interior point method. These optimization methods require price-demand gradients to exist. At a consumer level, however, many loads are often discrete. For example, electric water heaters and small air conditioner compressors operate on simple on/off signals, creating non-smooth load profiles. Second, most techniques, such as those in [1, 41, 53, 57, 66, 67], do not consider the network implications of implicit optimization methods used in their solutions. For example, if a Newton-Raphson based method [68] is used to converge around an optimal solution, all of the iterations are assumed to occur in the same time instant. When networked, this implies that the load profiles for consumers are stationary while a solution is found, eliminating the incorporation of time-varying information during convergence. While prior work has created the groundwork for feasible real-time transactive control at the consumer level [69], the behavior of most algorithms is undefined if communications are lost from some clients. The technique presented in this section is derivative-free, non-stationary, real-time, and resilient to a variety of network conditions.

Distribution scheduling presents a unique set of challenges that this section addresses. First, large numbers of low-budget consumers are interacting with the power market. This points to best-effort, existing network technologies for communication among market players that introduce jitter, congestion, and reliability issues into the communication infrastructure. The technique in this section provides mechanisms for overcoming unstable communications via continuity planning and efficient inclusion of outdated information where possible. Second, individual consumers may not have well defined cost information, and it may change with time. To optimize price in this context, derivative free techniques are used that make no assumptions about the structure of consumer's load profiles. Finally, since new consumer technologies such

as plug-in electric vehicles can respond rapidly to price signals, on-line price signal that changes rapidly without fixed windowing are used.

In order to address these challenges, the Nelder-Mead (NM) optimization heuristic [70] is modified to work on an online, non-stationary stochastic objective function. Standard NM is not designed to handle stochastic problems [71], so it is modified as to tolerate stochastic processes. The algorithm is further enhanced by accounting for non-stationary objective functions. The proposed optimization technique, though applicable to non-stationary stochastic problems, is still not robust to network failures as each iteration requires an objective function evaluation that requires network communication. The NM technique is further enhanced by providing network-immune function evaluation via state projection and estimation.

Using this technique, it is possible to provide online, dynamic pricing for an example distribution system's power market. The technique is able to increase efficiency by 64% over optimal 5-minute time-of-use pricing by rapidly issuing dynamic price values to consumers. If a rapid transient occurs, then the online, algorithm is able to reduce residual power by 85% over an a priori time-of-use approach. The performance of the technique is evaluated for a variety of network conditions, and it is shown that even under high latency situations, with 20% of clients disrupted or less, the solution still performs better than time-of-use solutions.

The rest of the section is structured as follows: Section 3.2.2 provides additional background on the DSM problem. Section 3.2.4 provides an overview of related work. Section 3.2.7 covers the solution technique, and Section 3.3 investigates networking impacts. Section 3.4 evaluates the performance of the algorithm presented in this section, and Sections 3.4.7 and 3.4.8 discusses and concludes it.

### 3.2.2 Background

#### DSM, Dynamic Pricing, and the Power Grid

In the power grid, the goal is to constantly match supply with demand [1]. The mismatch in supply and demand is known as *residual power (RP)*, and without widespread energy storage devices, it must be immediately corrected. Small amounts of residual power is absorbed by automatic generation control and system inertia. Large amounts of RP, however, can result in grid overload and cause brownout and blackout conditions. To minimize RP today, power utilities implement day-ahead scheduling based on energy usage forecasts. Market players submit bids some time in advance to meet the forecasted energy needs at various intervals spread throughout the day. While the vast majority of energy is scheduled a day in advance, a small amount of generation participates in a "real-time" market (RTM) that attempts to rectify the residual power resultant from forecasting errors and unexpected transients such as generator outages. In this case, real-time is relative to day-ahead, and the market for one real-time period may close well in advance of that period. The prices in the RTM can range from \$30 in one 5-minute window to \$300+ in the next. Dynamic pricing and demand side management are designed to reduce this variance by bringing more market players, especially consumers, into the real-time domain.

Demand side management (DSM) [65] entails a large suite of technologies that bring real-time flexibility to the smart grid. The components include digitally controllable loads, network protocols, metering devices, and other pieces required to adapt power supply and demand to a control signal, used by transactive control systems [72]. A general dynamic pricing objective function [73], used by DSM, is shown in Equation 3.15. It minimizes RP by controlling the locational marginal pricing (LMP),  $\lambda$ , throughout the system.

$$\arg \min_{\lambda} \sum_{n \in N} \left( \left| \sum_{i \in n} P_i(\lambda_n) + P_t(n) \right| \right) \quad (3.15)$$

where  $n \in N$  covers all the nodes in the power grid,  $P_i(\lambda_n)$  is the power used or produced (negative) by each client at price  $\lambda_n$ , and  $P_t(n)$  is the inter-nodal power transfer constrained by transmission capacity.

### 3.2.3 Real-Time Communication

The algorithmic core to dynamic pricing is a numerical optimization technique used to solve Equation 3.15. The fundamental basis for dynamic pricing is that the distributed consumers have access to changing, private information that is beneficially incorporated into the pricing optimization problem, i.e.  $P_i(\lambda_n)$  is actually  $P_i(\lambda_n, t)$ , a continuous-time (though likely discrete) function. This inherently requires constant communication between the DSM market players to adapt to changes as the system evolves.

As DSM methods solve for new prices, they are broadcast system-wide to the market players. The market players respond by adjusting consumption and production values that the DSM algorithm samples for its next price calculation. This process happens in two stages: negotiation and actuation. In the negotiation stage, the prices are hypothetical and used to improve solution quality, for example in a Newton-Raphson iterative process. Constraints or other dynamic information are exchanged at this level, depending on the solution algorithm used for DSM. The actuation stage sends contractually backed prices on which the consumers act. Figure 3.6 shows this process.

Many DSM techniques convert the continuous-time DSM problem into discrete windows for which a solution can be generated by a stationary optimization technique. Once the algorithm has converged, a new window can be created and the algorithm advanced. A practical problem with this approach is that new information cannot be incorporated into the problem until the next window occurs. Additionally, the functions  $P_i$  are owned by the distributed market players and must be syn-




Fig. 3.6. Communication pattern: Negotiation phase and actuation phase.

chronously frozen. Any deviation by the individual market players from the frozen  $P_i$  can invalidate convergence properties that rely on stationary functions.

### Nelder-Mead and Derivative Free Optimization

Many optimization techniques analytically derive  $P_i$  and use its formulation in an interior point method or gradient descent to solve Equation 3.15. These methods make assumptions about client behavior when formulating  $P'_i$ , and then iterative methods such as the Newton's method are used to optimally find  $\lambda$  as the roots of Equation 3.15. These techniques, however, rely on the derivability of  $P_i(\lambda)'$  and may not exist for  $P_i(\lambda, t)$  or at all for clients with discrete behaviors.

Derivative free optimization techniques such as the Nelder-Mead method [70, 71] solve the optimization problem via heuristic searching without relying on derivations of  $P_i$ . These techniques iteratively evaluate the objective function (RP) and make algorithmic decisions on which point to evaluate next based on prior observations. The convergence properties of these methods are not well guaranteed, but they do usually generate sound solutions to the dynamic pricing problem when used in this section.



## Practical Limitations to DSM

Optimizing DSM requires extensive network communications to occur over a wide area network. Geographic limitations restrict the speed at which messages can be exchanged, and throughput limitations restrict how many messages can be transmitted and received for any window of time. Therefore the cost of optimization is summarized as the number of function evaluations required to reach an optimal solution.

The computational requirements for large optimization problems solved by Nelder-Mead are rather minimal sorting and summation components across a fixed number of historical values. Computational restrictions are therefore ignored. The networking restrictions for a price-power exchange is roughly 100 bytes of data in each direction (64 bytes for a TCP datagram and 32 for a value). A 10-Gbit connection could therefore sustain a theoretical 1.5 million exchanges per second, a suitable amount for a large power market.

### 3.2.4 Related Work

### 3.2.5 Existing DSM Techniques

A wide range of existing DSM techniques provide mathematical and heuristic basis for optimizing power delivery in the smart grid. Work in [41] creates a general concave optimization problem around dynamic pricing and solves it using an iterative gradient projection method. This approach, however, utilizes gradients in consumer's utility and solves optimization problems using fixed windows of time. Work in [53] optimizes DSM with plug-in electric vehicle (PEV) consumers acting as the model for  $P_i$ . This technique also uses integer programming or interior point methods to solve for stationary dispatches. A common theme in DSM approaches [1, 57, 65] is to segment the energy market into windows in which the system is stationary and a solution can optimally converge before adjusting the market. Most methods rely on stationary problem parameters and ILP/MILP/IPM to solve the distributed optimization,

and they do not consider the communication patterns or overhead associated with distributed gradient and function evaluations.

Work in [74] provides a distributed method for optimizing prices in DSM. The technique, however, relies on broadcast messages to all connected clients which is not a scalable solution. Additionally, the problem is assumed stationary between iterations. In this section, a system that accommodates the delays due to real-time communication and does not require the system to be stationary between iterations is proposed.

### 3.2.6 Existing Latency Studies

There has been some work on DSM with lost communication messages in [57]. The authors frame the DSM problem using an interior point method to solve for the prices, like several other techniques. They then analyze the convergence on lost communications by solving the problem using stale, outdated information. Work in [56, 75] and the results (also derived in [57]) show that convergence is still possible when the gradient/lagrangian are stationary but delayed in time. In [56] the convergence rate is shown to slow as a function of the amount of stale information incorporated into the convergence process. Under latency conditions  $\tau$ , the solution convergence error for gradient descent is bounded by  $\sqrt{\tau/T}$  where  $T$  is the number of negotiation iterations [76]. These techniques, however, rely on stationary (and existing) gradients at the market players. In this section, a system that is based on derivative free optimization and does not require knowing the gradients at the market players is proposed.

### 3.2.7 Solution

#### Definitions and Assumptions

The market consists of players or clients whose only action is to consume or produce power. A central price-setting entity (assumed to be a distribution-level utility) dispatches prices to clients as either hypothetical (negotiation) or actionable as a control signal, and the power utilized by each client at a particular price point is assumed authentic. Consumer/producer constraints are manifested in the cost function as asymptotes. Non-participatory clients are modeled as offsets in the nodal power balance.

#### Overview

In the nodal power balance model, the consumers are a distributed set of market players connected to a central price coordinator. The coordinator controls the price signals at each node to balance power in the grid. To do this, the market players register with the coordinator and provide their real-time consumption levels in response to a hypothetical or actionable price that is periodically transmitted to the clients. The coordinator decides, via a pricing algorithm, subsequent prices in response to the power levels provided by the market players. Whenever a new price is determined to be more efficient than the existing price, it is made actionable. To determine optimal prices, the ISO periodically transmits prices and waits for a fixed period for a reply. In the event no reply is received from a market player, a substitute value for power is used. In the event information arrives after a period has expired, it may or may not be useful to the ISO for calculating future prices.

#### Non-Stationary Objective and Time-Scales

The underlying algorithm is designed around an implicitly non-stationary objective function. Each iteration of Algorithm 1 takes non-zero time, and this has an

influence on the behavior of  $f(x, t)$ , the objective function. A stationary  $x$  with a varying  $t$  will ebb and flow with the overall supply and demand of electric power in the grid, and sometimes during transient events such as faults, the change between  $f(x, t - \epsilon)$  and  $f(x, t + \epsilon)$ , denoted  $\Delta f()$ , can be dramatic. The standard Nelder-Mead algorithm on which Algorithm 1 is based will not converge if  $\Delta f()$  escapes the contracting simplex and the resulting finite, constrained search space. Several modifications are made to allow large  $\Delta f()$ .

### Nelder-Mead Optimization

Algorithm 1 contains the approach used to solve the optimal price in the system. The algorithm is based on a standard Nelder-Mead algorithm [71] but modified to support non-stationary problems. Three modifications are provided to enable the algorithm to perform online optimizations on non-stationary objective functions. First, the search space used by NM is modified to prevent collapse so that transients can be detected. Second, the algorithm is updated so that cached function values are updated to improve point-ranking. Finally, the algorithm is changed to enable quick retracing of the simplex space during its shrink operation.

In the algorithm listing, the function evaluations are explicitly demarcated to show where communication must occur as the evaluation of  $f(x)$ . Since the power consumption is determined by distributed clients, the  $x$  term must be transmitted to each client. The clients reply with  $f_c(x)$  and  $f(x) = \sum_{\forall c} f_c(x)$ .

**Simplex Collapse** The standard Nelder-Mead algorithm collapses a simplex around an optimal point in a search space. When the simplex is sufficiently small, or the standard deviation of  $f(X_1), \dots, f(X_n)$  is less than a threshold, the algorithm terminates with a solution (however, the online method never terminates since the objective function is non-stationary). If  $f(x)$  is non-stationary, then the optimal point may move outside of the simplex and outpace the reflection point's growth, especially if the simplex is small and the change in  $f$  is large. In the standard Nelder-Mead algo-

---

**Algorithm 1:** Modified Nelder-Mead

---

```

1 Initialize  $X_{1..n}$ ,  $f(X_{1..n})$ 
2 while true do
3   Sort  $X_{1..n}$  by  $f(X_{1..n})$  ascending  $f(X_1) \leq f(X_2)$ 
4   Set  $X_1$  as actionable
5    $X_o = \bar{X}_{1..n-1}$ 
6    $R_v = \text{rand}(-0.5, 0.5)$ 
7    $R = R_v(\omega + \beta/(\text{std}(X_{1..n}) + \beta))$ 
8    $X_r = X_o + \alpha(X_o - X_n) + R$ 
9   Evaluate  $f(X_r)$ 
10  if  $f(X_1) \leq f(X_r) < f(X_n)$  then
11     $X_n \leftarrow X_r; f(X_n) \leftarrow f(X_r)$ ; continue
12  end
13  if  $f(X_1) < \pi$  then
14    Re-Evaluate  $f(X_1)$ 
15  end
16  if  $f(X_r) < f(X_1)$  then
17     $X_e = X_o + \gamma(X_o + X_n)$ 
18    Evaluate  $f(X_e)$ 
19    if  $f(X_e) < f(X_1)$  then
20       $X_n \leftarrow X_e; f(X_n) \leftarrow f(X_e)$ ; continue
21    else
22       $X_n \leftarrow X_r; f(X_n) \leftarrow f(X_r)$ ; continue
23    end
24  end
25  Re-Evaluate  $f(X_{n-1})$ 
26  if  $f(X_r) \geq f(X_{n-1})$  then
27     $X_c = X_o + \rho(X_o - X_n)$ 
28    Evaluate  $f(X_c)$ 
29    if  $f(X_c) < f(X_n)$  then
30       $X_n \leftarrow X_c; f(X_n) \leftarrow f(X_c)$ ; continue
31    end
32  end
33   $X_b = X_1$ 
34  for  $k=1..n$  do
35     $X_k \leftarrow X_b + \sigma(X_k - X_b)$ 
36    Evaluate, Store  $f(X_k)$ 
37  end
38 end

```

---

rithm, the simplex can only expand in the case that  $f(X_r) < f(X_1)$ , the best point. This means that the algorithm will become stuck in practice, especially since  $f(X_r)$  will only become worse as the optimal point moves away from the simplex. The first modification, in lines 6, 7, 8, adds random noise, inversely proportional to the size of

the simplex, to the reflection point. This allows the simplex to expand, as a function of  $\omega, \beta$ , when the search space is very tight. If the reflection point is accepted, then the centroid  $X_o$  can move a substantial amount, accelerating the pace by which the simplex can track the non-stationary optimal point.

**Inaccurate Ranking** The sorting operation on line 7 operates on historical, memoized function evaluations, each of which takes a non-trivial amount of time to calculate. Consequently, if at any point  $X_1$  is close to the optimal  $X$ , then its rank will become difficult to change by subsequent evaluations unless the simplex shrinks. Line 13 contains a condition in a parameter  $\pi$  to re-evaluate the optimal point  $X_1$ . Since the residual power problem is convex with a optimal point at zero, the amount of re-evaluation requests can be throttled this parameter. Intuitively, when the objective value of  $f(X_1)$  is large, the probability that  $f(X_r) < f(X_1)$  will be large. This compounds with the simplex expansion in Section 38, so that when the optimal point leaves the simplex, the probability that the simplex updates with  $X_r$  increases.

**Reflection Retracing** The most expensive operation in terms of function evaluations is the, reduction/shrink part of the algorithm starting on line 34. In this phase, the entire simplex is reduced around the best point  $X_1$  as  $X_b$ . If the optimal point moves quickly enough away from the existing simplex, then the condition that the reflection point is worse than any existing point,  $f(X_r) \geq f(X_n)$ , will occur. Additionally none of the exploration points will be better than  $f(X_n)$ , by assumption. As a result, the algorithm will shift and re-evaluate every point in the history, including  $X_1$ . This expensive operation requires  $n$  function evaluations to complete, but it will enable the efficient movement of simplex provided  $\sigma$  is not too small. By also re-evaluating  $X_1$ , the ranking of the new points can be updated.




Fig. 3.7. Faults: When the communication channels are interrupted, the algorithm operation becomes blocked (left). With a recovery mechanism, a new price and power level can be used to continue operations (right).

### 3.3 Networking and DSM

The behavior of most DSM algorithms is undefined whenever information does not successfully arrive in each iteration, as shown in Figure 3.7. For example, Algorithm 1 will not complete with even a single straggler in the function evaluation phase. This section presents methodologies used to recover from non-linearity experienced in networked and distributed systems.

#### 3.3.1 Latency Influences in DSM

##### Negotiation

Traditional optimization techniques are designed to operate iteratively in lockstep. The prices for the next iteration depend on the power levels of the previous iteration, and these are assumed available at each iteration. In a large scale system, however, this means that the real-time iteration speed is bounded by the slowest market player. To solve this issue, latency can be accounted for by introducing stale information into the optimization problem when communication is interrupted. Optimization occurs

around  $z = f(x(t - \tau(t)))$  instead of  $f(x(t))$ , for example. Each iteration is assumed to take  $T_s$  seconds to complete before moving to the next step.

The relationship between  $\tau$  and  $t$  is maintained by wall-clock speed of each iteration,  $T_s$ . Each transmission of  $z(t)$  takes  $\tau_R(t)$  wall-clock seconds. Each overall iteration blocks for  $\tau_I$  seconds, waiting for all replies to arrive, then  $\tau(t) = \lfloor (\tau_R(t)/\tau_I) \rfloor$ . This maintains the relationship that error reduction is bounded by the real-time speed of the communication network. If the number of iterations completed in a window of time is doubled, and  $\min(\tau(t)) \geq 1$ , then  $\tau'(t) = 2\tau(t)$ ,  $\sqrt{\tau/T} = \sqrt{2\tau/2T} = \sqrt{\tau'/T'}$ .

The algorithm presented in Section 3.2.7 operates with a memoized history. Latency is captured by each function evaluation, where  $x$  is sent out to all the clients, and they reply with  $f_c(x)$ , and the objective function is  $\sum_{\forall c \in C} f_c(x)$ . Continuity is guaranteed by using an approximation function for  $f_c(x)$ . Initially all values for  $f(x)$  are populated with the projection function and then replaced as new values arrive via communication systems. In this case, if  $\tau_c(t) > 1$ , it will not be incorporated in the initial evaluation of  $f(x)$  in the algorithm. Instead, the historized value of  $f(x)$  is updated when the new information arrives. This allows for efficient a posteriori incorporation of stale information, manifested via re-ranking of  $f(X_{1..n})$ . The accuracy of the projection function will impact how well the method survives high latency situations.

## Actuation

Latency in the actuation loop for DSM introduces potential state inconsistencies in the system. There are three hypothetical states in the system: the a posteriori oracle-optimal  $x_O(t)$ , the DSM's intended state  $x(t)$  and the experienced state  $x(t - \tau(t))$ . The DSM is intending to minimize the objective function  $f(x)$  which captures the inefficiency of the state. At a base level, the penalty for actuation delays is  $f(x(t - \tau(t))) - f(x(t))$  while the degradation performance of the system is  $f(x(t - \tau(t))) - f(x_O(t))$ .






Fig. 3.8. Asynchronous state traversal from  $x(t_0)$  to  $x(t_1)$  via several potential paths, each a function of communication latency.

In existing DSM techniques, timescale for negotiating is much shorter than actuation, however, which leads to granularity issues in tracking the true impact of  $\tau$ . For example, the time-scale for actuation may be 5 minutes. Latencies between 0 and 5 minutes are captured by  $\tau = 0$  i.e. the actuation is based on a projection 5 minutes in the past, and as long as the message arrives by 5 minutes, there is no consequence at this level. The problem is complicated due to the lack of a distributed agreement protocol in the system. For example, consider a system in which  $\tau(t) \geq 1$  for all market players. Since information is received asynchronously, the evolution from  $x(t_0)$  to  $x(t_1)$  is asynchronous, as shown in Figure 3.8. With the right latency values, the asynchronous state traversal may transit state spaces that perform poorly. Ideally, the movement between states would be instantaneously, as modeled in DSM. However, the speed at which traversal occurs, a function of communication latency, impacts the performance of DSM. The movement between states is necessitated by the encroaching "degraded" operation space (i.e.  $S(x(t))$  is becoming worse over time). For this reason, costly distributed agreement protocols can have a high penalty.

The algorithm presented in this section, however, actuates whenever any new, more-optimal  $X_1$  value is found (i.e.  $X_1(t-1) \neq X_1(t)$ ). At this level, the system responds much more quickly to changes in the power grid. Degraded states are much more easily avoided. Stability can be ensured by coordinating, a-priori, the load points that the consumer will use if communication is disrupted.

### 3.4 Experimental Results

This section provides and analyzes two primary experiments. In the first Experiments, 1-3, the base optimization technique is evaluated and compared to the optimal performance of other techniques. This section also analyzes the algorithm's sensitivity to parameters described in Section 3.2.7. Experiment 4 analyzes the algorithm's ability to withstand uncertain network conditions expected to arise from best-effort networks.

#### 3.4.1 Experimental Setup

The experiment uses agent-based modeling to capture the behaviors of an example smart grid setup. There are three primary agents: the consumer/producer, the market manager, and the communication agents. The consumer/producer agent models a generator or consumer in the system, the market manager implements Algorithm 1, and the communication link accounts for latency and congestion behaviors in the system. A central event-based queue manages the interactions between the agents, and an external scenario file describes the transient behaviors of the agents.

#### Communications Model

The communication between the agents is managed via a FIFO queue. The delivery time for a message is calculated as latency plus transmission time (size/bandwidth) with a maximum queue size that introduces dropped packets. The bandwidth and

Table 3.1  
Example parameters used in Figure 3.9

	$P_{\min}$	$P_{\max}$	$\lambda_{\min}$	$\lambda_{\max}$
Consumer	0	110	0	100
Generator	-90	0	30	80

latency parameters can be controlled to introduce latency via agent-to-agent message congestion or via external congestion.

### 3.4.2 Load and Generation Model

This experiment models supply and demand as scaled sigmoid functions, as shown in Equations 3.16, 3.17 as  $P(\lambda)$ . Figure 3.9 shows two example curves with the corresponding residual power at each pricing point, based on the parameters in Table 3.1, where the generator power level is presented as positive but actually negative. The sigmoid function was chosen to allow responsiveness to price while also facilitating asymptotic behavior at the extremes, capturing market-enforced constraints.

$$\lambda_s = 6 * \frac{\lambda - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} - 3 \quad (3.16)$$

$$P = \frac{P_{\max} - P_{\min}}{1 + e^{\lambda_s}} + P_{\min} \quad (3.17)$$

For the experiments conducted in this section, unless specified, the parameters used are as follows. There are 2 generators  $G$ ,  $N_G = 2$ , and 100 consumers  $C$ ,  $N_C = 100$  with a target power  $T_P = 100$ . The mean  $P_{\min}$  for the generators is  $\frac{P_T}{N_G}$ . The  $P_{\max}$  for the consumers is  $\frac{3P_T}{2N_C}$ . The  $\lambda_{\max}$  for  $C$  is \$100, and the  $\lambda_{\min}$  is \$0. For the generator it is \$80 and \$30 respectively.

The non-stationary components are as follows. The generator and consumer have a time-constant  $\tau_G, \tau_C$  as  $\frac{3600s}{\pi}$  for a period of two hours. The magnitude of change  $A_G, A_C$  is 0 and  $\frac{0.35P_T}{N_C}$  respectively (generation is not time varying by default). For generators, only  $P_{\min}$  is modulated as  $A_G * \sin(\frac{t}{\tau_G})$ . For consumers, both  $P_{\max}$  and




Fig. 3.9. The supply (generator) and demand (consumer) are shown for varying market prices. The residual power shown is the imbalance between the generator's output and the consumer's input. Since residual power can disrupt the grid and lead to inefficient generation, the goal is to optimize the market price such that the residual power is minimized.

$P_{\min}$  are modulated by  $A_C * \sin(\frac{t}{\tau_C})$ . This creates periodic swings in demand that require constant updating of  $\lambda$  to maintain system optimality.

### Transient Events

To analyze the effectiveness of the online algorithm at handling unpredictable transients, step changes in demand are introduced into the model. To model step behaviors, the  $P_{\max}$  and  $P_{\min}$  parameters are shifted by  $\frac{P_T}{2N_C}$  and  $\frac{P_T}{10N_C}$  respectively. This shift is done at different points in time depending on the particular experiment.

#### 3.4.3 Experiment 1: Analysis under Normal Conditions

In this experiment, the DSM technique is subjected to a two market scenarios: one predictable and one unpredictable. The network conditions are assumed to be

favorable such that 60 iterations can occur per minute (1 per second). A time-changing scenario is run for one hour of simulation time and compared to an optimal windowed output function, time-of-use (TOU). The primary difference between the online method and the TOU method is the granularity of control.

In the first experiment, the market price is plotted across time for one hour, shown in Figure 3.10, under the time-varying parameters described in Section 3.4.2. Three different profiles are presented—the online, optimal, and 5-minute ideal time-of-use (TOU) values. The online solution closely matches the optimal with a lag for iteration-execution of one iteration (one second). The TOU solution is only optimal at one point per window in this scenario. It is worth noting that as the window size of TOU approaches zero, it matches the optimal solution (only with perfect prediction, since TOU is scheduled in advance).




Fig. 3.10. The  $\lambda$  price values are shown for a particular one-hour scenario. The optimal price is constantly changing, and the online method is attempting to match it. The time-of-use (TOU) method provides pricing windows with perfect prediction. The online method is much closer to the optimal price than the TOU approach due to increased granularity of control.

The optimality of the online approach in the residual power domain is shown in Figure 3.11. The online solution never perfectly matches the optimal solution, but it gets close at the apex of the price curve where the system is slowly changing. The TOU technique is only optimal at one location per window and suffers more when the slope of change is highest. The online method reduces residual power by 64% over the 5-minute time-of-use approach.




Fig. 3.11. The residual power in the market must be absorbed by automatic generation control. Less residual power indicates optimality in the pricing method. The online method is able to outperform the TOU method when changes occur in the system due to increased time-precision, reducing residual power by 64%. When the system is not changing rapidly, at the 30-minute mark, both approaches are comparable.

#### 3.4.4 Experiment 2: Unpredictable Step Change

This experiment adds an unpredictable step change to demand, as described in Section 3.4.2, to the prior experiment. At 20 minutes in, a change in demand occurs to simulate a large load suddenly coming online. At 40 minutes, the load is removed and the system resumes its predictable behavior. Figure 3.12 shows the price response

to this situation. TOU is unable to adapt to the transient since it is calculated a priori while the online method is able to observe and respond to the changing market condition.




Fig. 3.12. A step load is added and removed at 20 and 40 minutes respectively. The TOU method, calculated a priori, is unable to adapt to the change in the system. The online method lags the optimal method during the step transient, but it is able to adapt to the changing market conditions.

The residual power from this experiment is shown in Figure 3.13. The online algorithm adapts to the transient by observing market conditions and making real-time updates to the price. The algorithm rapidly improves in a few market iterations, resulting in a 85% reduction in residual power over TOU. Higher speed iterations (sub one-second) could enable outage avoidance by curbing demand quickly in a transient such as a transmission line being disabled.

### 3.4.5 Experiment 3: Sensitivity to $\beta$

A key parameter to the adaptability of Algorithm 1 is the  $\beta$  parameter which controls the minimum exploration capability of the algorithm (e.g. the minimum size of the simplex). Figure 3.14 shows how the system responds to various values of  $\beta$




Fig. 3.13. The residual power is shown for a step change in load. The online method responds to the transient to curb residual power and outperforms the a prior TOU method by 85%.

during a transient that lasts between time 200 s and 400 s. The lower values of  $\beta$  allows tighter simplex solutions to exist. This hurts adaptability to change, both time-varying and transient. The higher values allow quick changes to occur in the simplex search space, and as shown in the figure, the high  $\beta$  values are quicker to adapt to change. The lowest values are unable to adapt to the transient and become optimal again once the transient ends. In some sense, the  $\beta$  parameter tunes the responsiveness or long-term smoothness of the algorithm.

Figure 3.15 shows a spread of  $\beta$  values and the average impact on residual power for scenario described earlier. Higher values of  $\beta$  give more flexibility to the system and allow it to adapt to changing conditions. In low-transient situations, the value of  $\beta$  has limited impact relative to other terms in the algorithm. This is because when the system is tightly converged on a particular solution, explorations in new directions yield dead ends, so there is limited impact.






Fig. 3.14. Different  $\beta$  values for Algorithm 1 are plotted for a step change at time 200-400 seconds. The high values of  $\beta$  improve adaptability to transients while the low values resist drastic changes to the price.




Fig. 3.15. The average residual power in the market decreases with higher  $\beta$  values during the transient period. Without transients, the  $\beta$  parameter has little impact on performance.

#### 3.4.6 Experiment 4: Networking Impacts

For this experiment, the prior experiment is repeated with increasing latency. In order to operate under latency, the ISO utilizes a power consumption estimate for

clients that do not respond within one timestep. Equation 3.18 captures the model used for this experiment.  $P_e$  is the pre-transient power and  $\tau$  is the time since an update was received from the client. Parameter  $J_o = -0.2$  offsets the estimate and  $J_v = 0.01$  adds random noise to the offset with  $R$  being a sample from the normal distribution with mean and variance equal to one. The  $J_t = 50$ s parameter controls the time-dependent decay of the estimate accuracy. Using these parameters, the online algorithm is harmed by non-reply from clients.

$$P' = (P_e(1 + J_o) + P_e * J_v * R) * \frac{\tau}{J_t} \quad (3.18)$$

The networking model for this experiment is a star or hub and spoke topology. Each client has two communication agents, one upstream and one downstream, interfacing with the online algorithm. The latency values are assigned to each link as a constant offset and the bandwidth is set to 1 Gbit. A random subset of communication agents have latency added, controlled by the fraction parameter. At 50 seconds, the latency is added to the scenario. At 100 seconds, a transient occurs, and at 200 seconds the latency is removed from the scenario which terminates at 300 seconds.

Figure 3.16 shows the performance of the online algorithm with varying fractions of the communication agents experiencing high-latency situations. The fixed-price solution shown captures the behavior if the perfectly projected price is used, absent adaptation to the transient. When too many clients are disrupted, the online algorithm is penalized by the inaccuracies introduced by Equation 3.18. Even with these inaccuracies, the online algorithm is able to withstand 20% of clients being disrupted before performing more poorly than the fixed/windowed solution. A more accurate estimation function would improve performance.

### 3.4.7 Discussion

The algorithm presented in this section is designed to accurately respond to quick transients in the power grid by using price signals. Much optimization, however, takes




Fig. 3.16. This figure shows the performance of the system with a varying number of clients experiencing latency conditions. As more clients are exposed to latency, the system is unable to adapt well to transients. Small populations are easily accounted for by the online adaptability of the algorithm, but the larger populations place a high reliance on the load-estimation function.

place in longer term planning. For example, work in [43] optimizes consumption centrally by observing each loads needs and flexibility and developing a global operation plan. The approaches are complimentary, however, since long term plans are inherently not robust to changing market conditions. Ideally, the algorithm presented in this section would be used in conjunction with global planning to create an optimal, robust solution.

A limitation to this approach arises from developing ultra-precise control. Much of the price fluctuations at fast timescales are viewed as noise, especially if the system has substantial inertia e.g. synchronous motors as loads. For this reason, there is usually resistance in moving from 15-minute to 5-minute windows, etc., especially at a wholesale level. It is worth noting, however, that solid-state electronics such as those found in electric vehicles can rapidly respond to market conditions.

### 3.4.8 Conclusion

In this section, a technique for optimizing dynamic pricing signals for distribution level power markets was presented. A Nelder-Mead optimization technique was modified to support non-stationary, non-smooth price-demand curves for consumers in electric power markets. Network-aware enhancements were applied to the optimization technique, and it was used to optimize dynamic prices for an example market scenario. The technique was able to reduce residual power by 64% over 5-minute time-of-use methods during stable scenarios and by over 85% during large transients. In future work, the particular patterns of latency that are most disruptive to the optimization problem will be analyzed, and that information will be used to model strategic adversaries and cyber attacks on the grid.

## 3.5 Technical Market Conclusion

This chapter covered technical market operations—the mechanics and algorithms by which market clearing prices are found. The two methods presented were analyzed from a network perspective and shown to be resilient in some cases and unstable in others. This chapter motivates the need for further strategic adversarial analysis, and the next chapter combines this chapter and the previous one into a concrete strategy space for launching network attacks to disrupt power markets.

IMPROVING THE RESILIENCE OF  
CYBER-PHYSICAL SYSTEMS UNDER STRATEGIC ADVERSARIES

VOLUME 2

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Paul Wood

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

August 2016

Purdue University

West Lafayette, Indiana

## 4. MARKET MANIPULATION

In this chapter, the economic models from Chapters 2 and 3 are combined into a unified model for strategic adversaries in cyber-physical systems. The first section of this work has been published in COMSNETS '16 [77].

### 4.1 Introduction

Dynamic pricing markets in the smart grid (SG) [58, 65] enable the optimization of physical resource allocation. NYISO, the power system operator for the state of New York, could realize as much as \$400 million [64] in annual efficiency gains by leveraging wide-area real-time dynamic pricing systems. These gains, however, depend on consistent, reliable communication networks to facilitate control signal and measurement exchanges. The control signals comprise the price signals that the utility sends to the consumers and the measurement signals comprise the readings of the electricity usage at the consumers. Consumer-grade networks are often unreliable or congested at times, and they are highly susceptible to denial-of-service (DoS) attacks that disrupt communications entirely. As researchers pursue SG and other incentive-driven network control systems, they need tools to first understand what will be impact of outages of the network infrastructure on the demand-driven pricing mechanism on the SG and then how to mitigate the impact of this increasing attack surface and improve system resilience.

Electric power markets suffer from volatility because electricity is not easily stored. This volatility is expected to become much more acute with the increasing use of renewable energy sources, such as, solar and wind, that depend on the weather patterns. This creates a constant need to match supply with demand, but presently demand is inelastic and unaware of real-time market conditions. Supply and demand has histor-

ically been predictable which has limited this mismatch of supply and demand to a tolerable level. New renewable energy resources (RER) such as solar and wind driven supplies, however, have reduced this predictability to the point that technologies like roof-top residential solar are becoming cost-prohibitive [2] to integrate into the grid. The future smart grid is designed to bring elasticity to demand via techniques such as demand response (DR) [65] and transactive control (TC) [72] so that RER's can be better integrated and system efficiency improved. These techniques, however, rely on extensive communication infrastructures to coordinate wide-area energy consumption.

Transactive control enables distributed, independent control systems, operated by independent *actors* or market players, to coordinate via incentive-driven signals (prices). For example, the set point on an air conditioner may be sensitive to the cost of electricity in an automated way. This incentive can be set a priori via time-of-use pricing, but it is not sensitive to unpredictable changes in market conditions. Alternatively, a central market coordinator can negotiate with these automated loads by exchanging price and load information in real time with all of the actors. This negotiation process enables actors to rapidly respond to fluctuations in grid supply and by adjusting their energy usage based on price signals and their current exogenous needs. It is possible to disrupt the price signal negotiation, however, via attacks on the wide-area communication network. Since network attacks can influence the market price of energy directly, via control signal disruption, a *strategic adversary* (SA) can potentially launch attacks to manipulate prices in her favor.

When communications between market players are disrupted, the transactive control system becomes unable to influence consumption or production at those market players. For example if there is a spike in demand, the price signal should rise to curtail consumption and promote production. Producers who are aware of this signal increase their output and collect additional profits. An attack could disrupt the market signal at the producer, however, and as a result, power output would remain stagnant. Consequently, the market price may rise higher than it otherwise

would have to sustain equivalent demand curtailment, and this may benefit the other producers in the market. It could also lead to blackouts if the disruption is severe enough. If the SA is a producer, then direct financial benefit in the power market can be gained from such an attack. The SA can also benefit from the profits of multiple actors through various means such as investing in these actors. The ability for a network attack to benefit the SA is called the *attacker's incentive*, and this section focuses on measuring and reducing that incentive via defensive maneuvers.

Prior work in [78, 79] has shown that network attacks in smart grid control systems can disrupt price signals and provide benefits to subsets of consumers. These techniques, however, do not consider defensive maneuvers that the market players can use to protect themselves. Additionally, they rely on a strong adversary that compromises the entire market communication infrastructure. In this section, it is assumed that only the attacker can disrupt network links. Additional work in [20–24] has created a game-theoretic structure around attack and defense in control systems. These works do not consider the financial incentives of the attacker, however. Instead they focus on overall system performance or lower level dynamics and model the attacker as benefiting from system disruption rather than profiteering. In this work, game theoretic strategies for smart grids are combined into an attacker/defender game, with multiple defenders, that relies on financial incentives to motivate attack and defense.

This solution encompasses a method for estimating the attacker's incentive through attack strategies, mapping them to a game, and playing the game from a defender's perspective to minimize the attacker's incentive. First, a model for translating attacks and impacts on a smart grid is created to form a strategy space for the attacker. A dynamic market is implemented with communication links that can be disrupted via denial of service attacks to capture the attacker/defender strategies. From this space, the attacks are optimized to maximize the attacker's incentive (profits) by attacking communication links that distort the market to benefit the adversary. This is done via mixed integer linear programming (MILP). Then a model for defender is created that attempts to minimize the attacker's incentive by blocking certain attack strate-



gies via defensive investments (i.e. DDoS protection). Since the model has multiple defenders (actors), the impact of information sharing among the defenders on the reduction in the attacker’s incentive is also explored.

The solution is tested with a smart-grid based transactive control system. The baseline system optimizes power consumption by controlling the market price signal. A simulated communication network facilitates the exchange of price and load information. The attacker can choose which communication links to disrupt with a DoS attack, and the defender can choose some links to protect. It is shown that the baseline attacker incentives can be as high as 51% of overall operating profits. When the defender and adversary’s budget are equal, the attacker’s incentive is reduced by up to 70%. These results validate the utility of this section’s technique in optimizing defensive investments. It points the way forward for practitioners (such as, utilities) looking to deploy demand-driven pricing for electricity by showing how much resilience in the networking infrastructure is needed to assure a certain level of economic profit from the system.

The rest of the section is organized as follows. Section 4.2 covers the background in dynamic pricing markets and how they can be manipulated. Section 4.3 outlines the basic attack/defense strategy, and Section 4.4 expands the strategy to include information sharing among market players. Section 4.5 evaluates the strategies against an example dynamic pricing market, and the related work is discussed in Section 4.6. The section is concluded in Section 4.7.

## **4.2 Preliminaries**

### **4.2.1 Electric Power Grids**

Power grids are complex, interconnected systems composed of generators (sources) and loads (sinks). Each generator and load is connected to a series of transmission links (edges). A simple approximation of the energy system is a DC-load flow model which can be represented as a flow graph [4] where each asset (load, generator, edge)

in the physical system is an edge or node in the graph. Profit-seeking actors sell energy above cost (generators) or transform that energy into something more useful (loads). The system is most efficient when supply and demand are equalized since a surplus of power is dissipated as waste heat and a shortage of power causes brownouts, blackouts, and other grid stability issues. The imbalance of supply and demand is known as residual power (RP), and power grid operators strive to minimize this value. Dynamic market mechanisms [3] and demand response (DR) [65] minimize RP by either direct load control (DLC) or dynamic real-time markets. The work in this section focuses on power markets rather than DLC since the markets have a direct impact to profitability and thus attacker's incentive.

#### 4.2.2 Power Markets

Power markets utilize a variety of economic strategies to minimize RP (4.1) and maximize the system's social welfare (SW). The SW defines the global system benefit from energy transactions as shown in (4.2), where  $\omega_i$  is the value (consumers) or cost (producer) of power at each actor or market player  $i$ , and  $P_i$  is the amount of power consumed or produced by that market player. The parameter  $C$  penalizes the system for residual power with  $C \gg \bar{\omega}_i$ . While somewhat simple on the surface, the problem of maximizing SW is complicated by time-varying changes in  $\omega$  and the constraints on  $P$  that arise from power grid topologies and physical power constraints. To address these challenges, new smart grid models [3,53] allow real-time power markets to evolve with changing system conditions such as outages or unpredictable RERs.

The power market utilizes (4.3) to minimize RP, effectively maximizing SW. Each actor is exposed to the price  $\lambda$ , and they adjust their power output/input to optimize their individual economic situations. For consumers, their individual profit is  $SW_a = P_i \cdot (\omega_i - \lambda)$ . If  $\lambda > \omega$  for a consumer, then  $f(\lambda) = 0$  since the consumer would experience a net loss by consuming energy. Changes in energy needs or production




Fig. 4.1. The residual power (RP) and market price ( $\lambda$ ) are shown for an example grid scenario based on prior market solution work, later described in Section 4.5.1. The market experiences a demand surge in  $\bar{P}_i$  for consumers at  $t=100$  s followed by a reduction at  $t=250$  s. During the surge, residual power spikes until the market price is corrected.

are captured by (4.4). For example, a wind power producer has a very low  $\omega$  since wind is free and are thus driven by constraints in  $P$ .

$$\text{RP} = \sum_{\forall i} P_i(t) \quad (4.1)$$

$$\text{SW}(t) = \sum_{\forall i} \omega_i(t) P_i(t) - C \cdot |\text{RP}(t)| \quad (4.2)$$

$$P_i(t) = f_i(\lambda(t)) \quad (4.3)$$

$$\underline{P_i(t)} < P_i(t) < \overline{P_i(t)} \quad (4.4)$$

Online power markets optimize (4.2) by repetitively sampling  $P_i$  and updating  $\lambda$ . Each consumer receives a message containing  $\lambda$  and replies with  $P_i = f(\lambda)$ . Fig. 4.1 shows how a market evolves during a step transient in (4.4). The implementation of these systems, however, exposes security vulnerabilities that can be utilized by strategic adversaries to extract profit from the system. For example, the value of  $f(\lambda)$  may be based on an outdated  $\lambda$  during a communication outage thus reducing the SW. The defensive strategies presented in this work curtail impacts to SW in a cost-effective manner.

### 4.2.3 Profit Manipulation

Network disruptions have a direct impact on market price ( $\lambda$ ). Whenever a disruption occurs, the market player enters a zero-order hold mode. For market players, the price  $\lambda$  is fixed while  $P$  may change based on time-varying constraints. Consequently, the market loses its influence on power usage for a subset of market players whenever the network is disrupted. Fig. 4.2 demonstrates how the profits of a generator can be influenced by attacks on its communication link during the scenario shown in Fig. 4.1 and detailed in Section 4.5.1. The generator loses money if its own link is disrupted and can gain additional profits when some competitors' links are disrupted. Market players are retroactively charged the actual  $\lambda$  market price to promote market participation—otherwise self-disconnection would be a valid strategy.

## 4.3 Attack and Defense Strategy

Definitions:

$A$  set of market players

$I$  set of target network links

SW social welfare or profitability of the system




Fig. 4.2. The profitability for a generator is shown for two attack scenarios. At  $t=50$  s a DoS attack is launched on the communication link connecting the generator to the market (self-attacked) or another market player (others attacked), and it lasts until  $t=200$  s.

$SW_a$  profitability of market player  $a$

$IM[a, i]$  impact or change in profit realized by market player  $a$  when network link  $i$  is attacked

$P_i^{atk}$  probability of network link  $i$  being attacked

$P[a, i]$  probability of network link  $i$  being attacked, as estimated by market player  $a$

$D_i$  boolean indicating if network link  $i$  is defended

$C_i$  cost to defend asset  $i$

$A_i$  boolean indicating if network link  $i$  is attacked by the SA

$A[i, n]$  two dimensional  $A_i$  for  $n$  iterations in a game with imperfect information at the various actors

$O_a, O_i$  set of network links owned by market player  $a$ , owner of asset  $i$

$\sigma_a$  knowledge level of market player  $a$

#### 4.3.1 Players and Definitions

The set of market players  $A$  in the power market wish to maximize their profit  $SW_a$ , as defined in Section 4.2.2. The power system is comprised of a set of assets and their communication links in  $I$ . The term asset refers to both the physical system consuming or producing energy and it's associated target, the network link. There is a one-to-one mapping of assets (and thus targets) to market players defined as *ownership* such that one market player may own multiple assets. Each actor has a defensive decision to make for each asset that it owns— whether or not to invest in its defense  $D_i \in 0, 1$ . This decision has a cost of defense  $C_i$ . If the asset is attacked,  $A_i \in 0, 1$ , then the system experiences the impact IM from the attack, unless  $D_i = 1$  in which case the attack is assumed to fail via perfect defense.

#### 4.3.2 Attacker's Incentive

The strategic adversary (SA) attempts to profit from the manipulations described in Section 4.2.3 by launching network attacks on the links that interconnect the market players with the market mechanism and the price signal  $\lambda$ . Each attack results in a change in the profitability of each market player, and this is captured in the impact matrix  $IM[a, i]$  [4]. In this section,  $IM[a, i]$  is estimated via dynamic market simulations (Section 4.5.1) by approximating the market conditions for each player and evaluating the resulting changes in profit in the market. For example, the impact of attacking each network link on the generator in Fig. 4.2 is summarized by IM. The SA wishes to maximize the gain in profit for some market players with whom she has a financial interest, shown in (4.5).  $A, I$  is the set of actors  $A$  and network links  $I$  to attack and profit from as the attacker's strategy.

$$\operatorname{argmax}_{A,I} \sum_{a \in A, i \in I} \operatorname{IM}[a, i] \quad (4.5)$$

The probability of an attack on target  $i$  is proportional to the attacker's incentive gained from that attack. Abstractly, the target  $i$  could be any perturbation in the system—network outages, power plant disruptions, transmission line faults, etc. In this section, however, the focus is on a dynamic pricing system for the smart grid, and the targets are limited to network link disruptions. Similarly, the actors that benefit from the attack  $A$  are collections of consumers and/or generators participating in the dynamic market.

### 4.3.3 Defensive Maneuvers

The market players in the system can estimate the IM via their own impact analysis. Using their individual IM, they can also estimate the attacker's strategy and use it to construct a corresponding defensive strategy. Without any budgetary constraints, the defenders will protect all the targets in  $I$  by investing in high capacity, secured network links. Budgets are limited, however, so defenders must optimally select targets to defend. Section 4.4 describes how the defenders can have different views on the system parameters and still coordinate a defense.

## Underlying Game

The impact matrix IM is computed by assessing the *underlying game*, i.e. the power market, with successful attacks, as described in Section 4.2.3. Two versions of the system are compared—in one version, the attack was successful and in the other no attack is present. The resulting change in profitability for each actor is summarised by IM as the difference between the profits for each market player in the two scenarios.

Each attack on the system causes an overall net-negative impact on profitability. The system operates at a global-optimal whenever communications are uninterrupted.

Table 4.1  
Example Impact Matrix (IM)

	T1	T2	T3
A1	-2	-2	3
A2	4	-4	-2
A3	-4	2	-4

Any perturbations that disrupt communications result in decreased efficiency because of suboptimal responses to market prices ( $\lambda$ ). Therefore, the sum across all actors for any given target is always zero or negative. Some actors, however, may benefit from competitor elimination, which is the basis for the strategic adversary's profit model. Table 4.1 shows an example impact matrix for three market players and three targets.  $A_i$  owns target  $T_i$ .

#### 4.3.4 Defensive System Overview

The defensive investment optimization problem is designed to minimize the attacker's incentive thus reducing the probability of attack and denying profits to the adversary (resource exhaustion). An impact model is analyzed for each target and assessed as an impact to the profitability of each market player (IM). Once the matrix is calculated, it can be analyzed strategically to determine the best defensive action for each market player as in [4]. Fig. 4.3 shows the system layout.

#### Defensive Investments

Each market player in the system has a choice to defend self-owned targets from attacks at a cost  $C_i$ . If this cost is less than the expected reduction in profits, then it is in the actor's best interest to invest in defensive measures. The expected impact is  $IM[a, i]P[a, i]$  where  $P[a, i]$  is actor  $a$ 's expectation that asset  $i$  will be attacked, based on the SA's optimal strategy. In game theory terms, this section follows the






Fig. 4.3. In the overall system flow, set of defenders and a strategic adversary each have a view of the system and its market. The system is exposed to a physical scenario and analyzed for a set of potential targets via a market mechanism simulation. From this simulation, the profitability of each market player is captured in a set of impact matrices (IM). Each market player has an independently calculated IM and thus a different defense strategy that can be rectified via collaboration.

Stackelberg model where the different parties move one after the other. In this case, the attacker's move is estimated and the defenders decide to defend the assets appropriately. The attacker does not have the ability to come up with a repeat attack after observing the defensive actions, so the Nash equilibrium point is not analyzed in this model. Due to the slow-moving nature of defensive investments, the defender's strategy is not immediately observable by the adversary.

## 4.4 Multiple Knowledge Levels

### 4.4.1 Multiple Underlying Games

Each actor  $a$  and the strategic adversary have their own underlying game  $G_a, G_{SA}$ , respectively, from which  $IM$  is calculated. This arises because the parameters  $\omega_i, f_i, \underline{P}_i, \overline{P}_i$  from Section 4.2.2 must be estimated by actors who do not own those assets, including the strategic adversary. Each actor assesses the impacts from their independent viewpoints of the system. The game  $G_a$  is derived from the ground truth game  $GA$  by adding noise to the above parameters. In the perfect knowledge model, all of the games are identical,  $G_{a_1} = G_{a_2} \forall a_1, a_2$ . Imperfect information is modeled by allowing the underlying games to diverge by sampling i.e. fictitious play.

The game  $G$  itself contains a set of dynamic parameters  $x$  ( $\omega_i, f_i, \underline{P}_i, \overline{P}_i$ ) that are used to determine optimal market price. Fixed components of the game are the ownership and the network structure of the system. Each market player wants to keep its parameters secret to maintain a competitive edge in the marketplace. The dynamic parameters, however, can be estimated by observing market conditions and surveying physical equipment infrastructures. Each market player therefore can establish a "noisy" view of the underlying game, as defined by (4.6). The dynamic parameters are sampled from a normal distribution of the ground truth game. Sign changes are not allowed because it is assumed that each market player knows if an asset is a producer or consumer. The parameter  $\sigma_a$  defines the knowledge level of the actor  $a$ , and it is applied to all parameters except parameters for assets that the actor itself owns  $O_a$ . Intuitively, this parameter models the amount of information shared among each other. Greater is  $\sigma_a$ , less is the information that actor  $a$  has.

$$\begin{aligned}
 x' &= \mathcal{N}(x, \sigma_a^2) \forall x \in GA, x \notin O_a, \\
 x' &= x \quad \forall x \in GA, x \in O_a \\
 x' &\in (-\infty, 0] \text{ if } x' < 0, \text{ else } x' \in [0, \infty)
 \end{aligned} \tag{4.6}$$

#### 4.4.2 Perfect Information Game

In the perfect information game, the strategic adversary and all actors have a perfect view of the system,  $G_{SA} = GA$ ,  $G_a = GA \forall a \in A$ . In this form of the model, there is a single, optimal outcome for the defenders. Since  $P[a_1, i] = P[a_2, i] \forall a_1, a_2 \in A$ , the defensive decision is the same for each actor, and if the costs are correctly distributed among the defenders, then there is a single globally optimal defense strategy. The maximization problem (4.7) is solved by the defenders via mixed integer linear programming (MILP). The maximum value of this equation is zero because if no target is attacked, then no defense is necessary. Practically, protecting a network link (e.g. via DDoS protection) has some cost  $C_i$  for establishing a more reliable communication channel. The defender that owns each link, must decide to invest in its protection or not based on the likelihood of attack and the financial impact of the link outage.

$$\max \sum_{\forall a \in A} \sum_{i \in I} A_i \text{IM}[a, i](1 - D_i) - D_i C_i \quad (4.7)$$

The strategic adversary, the driving force behind  $A_i$ , is playing a similar game in (4.5). Since everyone shares the same knowledge, the perceived impact at each market player is the same, and all actors agree on which targets should be defended. Both the attacker and defender may have constraints on  $\sum A_i$  and  $\sum D_i C_i$  due to budget constraints on how many assets can be attacked and defended, respectively.

#### 4.4.3 Imperfect Attack Strategies

The adversary is assumed to be perfectly rational (no anarchy) but may not have perfect knowledge of the system and subsequently makes suboptimal decisions. To capture this,  $A_i$  is evolved into a mixed probability-based strategy across several underlying games for the adversary. The SA has a single, optimal (pure) strategy per (4.5), and a mixed strategy is created by combining multiple pure strategies into a single mixed strategy. Multiple  $\text{IM}'$  are calculated for the strategic adversary's

underlying game  $G_{SA}$  that are derived from  $GA$  as define in (4.6), with the caveat that the SA owns no assets. Equation (4.5) is optimized for each  $IM'$  across  $N$  fictitious games, each with a knowledge level  $\sigma$ , as a noise ratio. This results in  $N$  strategies for each asset  $i$  as  $A[i, n]$ . Equation (4.8) is the calculation for the probability of attack on target  $i$  given the  $N$  fictitious games for the adversary. The outcome  $P_i^a$  is an average of the boolean strategies for each of the SA's hypothetical games.

$$P_i^a = \frac{\sum_{n \in N} A[i, n]}{N} \quad (4.8)$$

Defense with Mixed Attack Strategies: The defenders strategy, as captured in (4.7), is modified below in (4.9) to account for the fact that the SA may have a non-boolean attack plan. Previously,  $A_i$  was binary and now  $P_i^a$  is a rational number so that the defender is operating on a mixed strategy.

$$\max \sum_{\forall a \in A} \sum_{i \in I} P_i^a IM[a, i](1 - D_i) - D_i C_i \quad (4.9)$$

#### 4.4.4 Multiple Defender Optimization

The maximization problems presented earlier for optimizing defensive investments do not consider the scenario where multiple defenders do not have the same information level and are optimizing around different underlying games. Each defender's underlying game,  $G_a$ , is used in place of  $GA$  to calculate a mixed attacker strategy using (4.8). Each actor then has a different threat model  $P[a, i]$  based on  $G_a$  instead of  $GA$ . (4.10) is performed by each actor to complete the optimization of  $D_i$ . Only the owner of asset  $i$  can determine the value of  $D_i$ . This approach enables no single actor to have a global view of the system which accurately models how a large interdependent system would operate.

$$\max \sum_{\forall a \in A} \sum_{\forall i \in O_a} P[a, i] IM[a, i](1 - D_i) - D_i C_i \quad (4.10)$$

Cost Collaboration: This problem is supplemented with a collaboration method. The cost of defense of target  $i$  is proportionally shared among benefiting actors. Since defensive decisions are segmented by asset owner, and attacks against owned-assets are always damaging, there is no Price of Anarchy (PoA) in this defensive model.

## 4.5 Experimentation

### 4.5.1 Experimental Setup

The underlying game, as described in Section 4.2.2, is solved via an online Nelder-Meade (NM) [71] optimization technique. Each iteration of NM is assumed to take one second and requires one round-trip communication of  $\lambda$  and  $P_i$ . The model for  $f_i$  is given in (3.17). In the case of a consumer,  $\omega_i = P_{\max}$ . For a producer,  $\omega_i = P_{\min}$ . The source code and corresponding market model details are available at [80].

The model has 20 generators with an average  $P_{\max} = 0, P_{\min} = 5, \lambda_{\min} = 30, \lambda_{\max} = 80$ , and there are 100 consumers with  $P_{\max} = 1.5, P_{\min} = 0.3, \lambda_{\min} = 0, \lambda_{\max} = 100$  for a total of 120 market players. The consumers  $P_{\min}$  is modified as  $P'_{\min}(t) = P_{\min} + 0.30 \cdot \sin(\frac{\pi \cdot t}{3600})$ , and all the other parameters are agnostic to time. At  $t=50s$ , the targeted network links are disrupted such that the  $\lambda$  term is fixed for those assets. At  $t=200s$ , the links are restored and communication is resumed. At  $t=100s$  a step load is introduced by setting  $P_{\max} = 2, P_{\min} = 0.4$  for all consumers. The parameters are restored to the default values at  $t=250s$ . The simulation is executed for 400 seconds. This scenario can be seen in Figs. 4.1 and 4.2.

## Communication Topology

In the experimental model, the communication paths between the market organizer (NM algorithm) and the individual market players are independent. In practice [81], however, there will be interdependence between communication failures across the different market players as many of them will share common links at some

point in the communication path. For this reason, the 120 market players are distributed on a tree topology with 4 top tier network links, 12 mid-tier links (3 for each top tier link), and 120 leaf links to better capture the interdependent networking impacts on smart grid topologies as shown in Fig. 4.4. Since future communication topologies have not yet been determined, and because dynamic markets may not operate on the same infrastructure as existing smart metering technologies, the topology used is purely speculative. As concrete topologies evolve, they can be substituted into this experimentation framework to identify changes in strategy and crucial network links.




Fig. 4.4. The simulated network has four top-tier links, twelve mid-tier links, and one hundred twenty leaf links.

#### 4.5.2 Experiment 1: The Attacker's Incentive

In this experiment, the strategic adversary attempts to maximize her incentive, as described in Section 4.3.2. The strategy space  $I$  is the selection of links in the communication topology to disrupt. Each disruption results in a particular IM that is used to calculate the maximum attacker's incentive. Fig. 4.5 shows the attacker's

incentive as a function of the number of links that she can attack simultaneously. As the number of simultaneously disrupt-able links increases, the attacker's incentive also increases. The attacker's incentive plateaus, however, when the overall system performance degradation becomes the dominating factor due to large numbers of link outages. The AI plateaus because the system as a whole becomes less profitable whenever most network links are disrupted (e.g. the top-level link attacks).




Fig. 4.5. The strategic adversary attempts to maximize her incentive by disrupting network links. In the graph on the left, the adversary is disrupting leaf-links in the communication topology. In the middle graph, the adversary is disrupting mid-tier links, and in the graph on the right, the top tier links are disrupted. In each case, the attacker's incentive is maximized for the given targets attacked. The strategy shown is maximized from (4.5) and compared to the mean of a random target selection.

### 4.5.3 Experiment 2: Collaborating Defenders

In this experiment, the defense strategy presented in Section 4.4 is evaluated. The defenders attempt to reduce the attacker's incentive shown in Fig. 4.5 by securing particular network links, thus eliminating them from the attacker's profit pool. The market players at each mid-tier communication hub are joined together so that there are 12 owners with 10 assets each. Collaboration is then possible on the mid-tier and top-tier network links, and they are the focus of this experiment.

Fig. 4.6 shows the reduction in attacker's incentive for different target budgets and a fixed  $\sigma = 0.1$  for each owner. The defensive budget is progressively reduced relative to the number of attacked links. The defenders are able to significantly reduce the attacker's incentive in most large-attack cases at the leaf links.




Fig. 4.6. The attacker's incentive is reduced by defensive investments in communication links. As the number of links attacked increases, the number of links defended also increases. The effectiveness of the defense, however, is reduced by imperfect knowledge levels among the defenders ( $\sigma = 0.1$ ). Each line represents a different amount of aggregate defense budget, relative to the number of links attacked.

Fig. 4.7 shows the reduction in AI for a range of knowledge levels across the defenders. In this case, 75 links are attacked and 75 links can be defended. The AI is maximally reduced when the owners knowledge levels are maximized ( $\sigma \rightarrow 0$ ) indicating that collaboration can improve overall defense effectiveness.






Fig. 4.7. This figure shows the cumulative reduction in attacker's incentive for different defender knowledge levels when 75 targets are attacked and defended. Decreased knowledge levels (high  $\sigma$  values) results in ineffective defense. As defenders are unwilling and unable to collaborate on defensive investments, the system suffers overall from poor defensive strategies.

#### 4.6 Related Work

Game theory applications for the smart grid [20] have become an increasingly important component of power system optimization. The core goals of these games, along with other approaches such as dynamic market mechanisms [3], are to improve the social welfare of the smart grid by utilizing market forces to balance demand with supply. The usefulness of these games, however, has not been well studied in the context of a strategic adversary that seeks to maliciously profit from the system by launching attacks.

A separate but related set of game theories optimize the security of information systems [26, 82] by playing attacker/defender games designed to create a defensive strategy that is optimal for a given adversarial model. Most of these approaches,

however, utilize qualitative metrics for target valuation, costs, and attack success since it is difficult to value computer system breaches. The work presented in this section applies attacker/defender concepts to a concrete smart grid cyber-physical system in which the utilities of attack and defense are derived from their actual operational influences. This also allows information exchanges to have quantitative impacts on success metrics.

Several attacker/defender games or security games have been constructed around Stackelberg games [83] for solving defensive investment optimization problems and scheduling patrols [84]. These games solve an attacker/defender model where the defender moves first in response to a perceived adversary and have been extended to support multiple human-modeled adversaries [85] in a computationally efficient way. These models, however, do not address defenders who exist in a competitive environment. The work presented in this section analyzes attacker/defender games in a competitive environment.

The long-term financial impacts of attacks have been studied in [30]. Adversaries are modeled in [30] as having budgets that deteriorate with unsuccessful attacks, resulting in reduced attack viability. The model captures some of the economic factors in this section, but it does not make a connection between the physical system's behavior and the resulting financial outcome of attacks. Similarly, game theoretic techniques in [20–24] have proposed methods for determine how adversaries might manipulate the physical control systems via attacks, but they do not draw the financial connection between physical perturbations and adversarial profits. The work presented in this section focuses on the financial motivations of attackers and defenders, resultant from system perturbations, as an attack and defense planning tool.

## Market Model Attacks

Prior work in [78] has focused on data integrity attacks in real-time pricing in smart grids. This work provides a model for a strategic adversary manipulating pricing information in a real-time market similar to DMM. However, their work does not involve large-scale systems or ones in which loads are time-varying, utilities are non-linear, and does not focus on economic dispatch. This work improves on this state of the art by evaluating large scale systems with real-time operating constraints (as emulated in DETER), specifically focusing on the delay-type attacks.

Work in [79] has analyzed market responses to integrity attacks in smart grids. The work provides an adversary model and analysis for a game-theoretic model of dynamic pricing systems in smart grids. This work focuses on the real-time implications of information-flow failures rather than continuously manipulated price signals which is more consistent with resilience and broadly available attack vectors (DoS) as opposed to manipulation attacks which are harder to implement in practice. This work is improved by executing the demand-response market algorithms in real-time and establishing interrupt-able communication between the market players.

Resilience improvements to control algorithms have been considered in [86], however this work focuses on grid stability and automatic generation control rather than market stability or cases where communication occurs over wide areas.

## 4.7 Conclusion

In this section, a modeling technique was presented to connect the networking components of a dynamic pricing market with a security strategy to defend against a profit-motivated adversary. This model allows the economics of cyberattacks on power markets to be captured and used to assess the risk to assets in the system. The amount of information that competing market players share about assets in the system is also modeled and used to analyze the benefits of collaboration in a defensive context. Techniques were then applied to mitigate the attacker's incentive thus improving

system resilience. It was shown that the baseline attacker incentives can be as high as 51% of baseline operating profits. When the defender and adversary's budget are equal, the attacker's incentive is reduced by up to 70%. These results validate the utility of this section's technique in optimizing defensive investments. The model presented in the section and the simulation results show promising approaches to countering the growing threat of cyberattacks in smart grids.

In future work, models of online learning aspects of dynamic pricing markets could improve attack strategies. In such a model, the attacker seeks to learn through iterative attacks, which also reveal more information about the system and the defensive strategies. Conversely, the defender also seeks to learn of the attack strategy through a multi-round strategy. The next chapter covers one angle of this approach.

## 5. NETWORK ATTACKS IN MARKETS

This chapter introduces a real-time heuristic to attacking real-time pricing systems. The previous chapter covered a strategy that is solved via mixed integer linear programming (MILP), but the problem size becomes intractable as the number of elements in play increases. It also does not consider multiple time-instants. To overcome this limitation, this chapter introduces a real-time heuristic to attack (and defend) real-time pricing systems.

### 5.1 Introduction

In the emerging smart power grid, new control methods such as demand response (DR) and real time pricing (RTP) are under development to improve the efficiency and reliability of the power grid [65]. RTP methods utilize economic incentives across networked control systems to stabilize imbalances of power in grid supplies and loads, minimizing waste and maximizing renewable power integration. The growth of these systems counters the uncertainty of renewable energy resource outputs and allows power fluctuations to be absorbed by flexible loads. Much like networked control systems (NCS), RTP methods rely on wide area communication networks to send economic incentive signals to the flexible end load points. Since the time period of the power fluctuations can be small, of the order of minutes, it is important to provide such signals in a timely manner to the end points. However, strategic adversaries can disrupt these signals, changing them, delaying them, or dropping them altogether. A recent NESCOR report [87] cites blocked DR messages as highly ranked failure scenario ("DR.1"), and it has been shown by prior work [88,89] that disruptions of price signals can cause disruption to power consumers. This section further demonstrates

how a strategic adversary can orchestrate and profit from such disruptions in these real time pricing systems, along with a mitigation technique.

RTP systems are particularly susceptible to network attacks because of the tightly coupled price–power feedback within such systems. Much like in other NCS’s, an RTP controller samples a process variable, in this case the flows of electric power, and modulates a control variable to minimize power imbalance over a communication network. Unlike networked control systems, however, RTP systems must coordinate among devices that are owned by different economic entities, requiring pricing or incentive based control instead of direct modulation of loads and supplies. In price based systems, the amount of energy input and output at each end point is a function of market price, and the system modulates the price to establish an efficient market clearing such that the sum of power inputs and outputs in the grid is zero. Since the price and power signals traverse geographically diverse end points, it becomes feasible for a strategic adversary to disrupt the power grid by disrupting the communication channels on which the smart grid relies.

This section develops a strategic adversary who capitalizes on the arbitrage of prices caused by disruptions in the RTP communication network. Each disruption removes controllable loads and supplies from the RTP system such that the market price trends higher or lower than in the attack-free case. The attacker then picks a set of network targets and disruption times with a goal of maximizing arbitrage opportunity. For example, if an attack is predicted to raise the market price by an additional \$20, an adversary with a large-sized battery can buy power to charge the battery, launch the attack, and then sell it back for an additional profit.

First an adversary with only the ability to delay RTP communications is explored. The market price changes due to both traditional load swings and network attacks in real time, and the adversary makes market observations from an end user perspective in order to plan attacks. She then launches denial of service (DoS) attacks against other end users by using a parametric algorithm triggered on the real-time price signal. Then allow more complex, RTP price signal manipulation (e.g. man-in-the-middle

attacks) is allowed for comparison with the simpler DoS attack strategies. Using these strategies, it is demonstrated how an adversary can increase arbitrage profitability by 69% using delay attacks and 98% using signal manipulation.

A novel defense mechanism is then put forward that randomizes the adversary's view of the network targets for the load end points. Consequently, the adversary's plan of targeting certain end points calculated through its algorithm can no longer be faithfully executed. The difference between the adversary's view and the actual system can be controlled by a defense parameter, namely, how many end point network addresses to randomize. This defense mechanism can be deployed through moving target defense mechanisms implemented by the RTP system operator. The network address assigned to the end load points is permuted either periodically, based on market fluctuations, or based on indications of attacks to minimize the impact of network disruptions on the RTP signal.

This work builds on two closely related prior approaches in [78] and [88]. Prior work in [78] measured the impact of delay and integrity attacks on RTP feedback signals. The authors showed that carefully planned attacks can create large oscillations and instability in market price, potentially crashing the market. In this chapter, it is shown that profit-driven attacks can be successful without destabilizing the grid control loops or damaging grid equipment. The attacks that are studied here are more feasible to be launched and are more likely to stay under the radar, thus having the potential for greater impact. Further work in [88] analyzed the impact of arbitrary delay and modification attacks on the incentive signals in the RTP system. They modeled a strategic adversary that could manipulate an RTP system if both the incentive and load signals are known for each consumer in a game theory structure. They did not, however, consider attacks that evolve in real time—instead they focused on day-ahead market planning techniques. In this chapter, there is a focus on real-time attacks where the adversary is aware only of local market price. This defense mechanism also distinguishes us from the two previous closely related approaches which focus on the attack modeling.

In the experiments, it is shown that an adversary could potentially profit from an RTP system with a simple rechargeable battery and access to a denial of service as a service provider. The adversary is able to extract up to \$119 per day from the RTP market using DoS methods, 69% higher than without attacks. If the adversary is able to compromise 20% of the devices, revenue could be increased by 98% or more. In the future, if fluctuations increase and battery prices decline, the profit amount is expected to increase further. Then it is shown that a defender, utilizing the shuffling and deception strategies, can reduce the adversary's profitability by 30%.

In this chapter, the following novel contributions are made:

- A strategy for a strategic adversary to illicitly profit from a real-time pricing mechanism in the smart grid is presented. The attack relies on delaying communication on a subset of the network links and for a subset of time, given by the adversary's algorithm.
- A defense strategy customized for protecting against such market manipulation attacks is presented. The defense strategy can be customized to fit within a certain defense budget and the benefits scale proportionally to the defense cost incurred.
- The cost to launch an attack of the type presented here, the economic advantage that can accrue to the attacker, and the cost of defense, all based on real-world scenarios and data is quantified.

The rest of the chapter is organized as follows. Section 5.2 covers the market mechanism background for real-time pricing. Section 5.3 details the attacker's strategy in the RTP system. Section 5.4 details some defensive techniques to stop attacks presented in this chapter. Section 5.5 details the experimental setup, including the particular RTP system in use. Section 5.6 has the experimental results. Section 5.7 provides some topic discussion, and Section 5.8 concludes the chapter.



## 5.2 Background and Market Operations

The goal of RTP systems is to constantly match supply with demand. The mismatch is known as *residual power (RP)*. Without widespread energy storage devices, it must be immediately corrected since positive RP (power surplus) is shunted and wasted while negative RP (power shortage) causes frequency droop, brownouts, and possible equipment damage.

### 5.2.1 Real-time Markets

To minimize RP today, power utilities implement day-ahead scheduling based on energy usage forecasts. The vast majority of energy is scheduled a day in advance (e.g. by time of use (ToU) contract), but a small amount of generation participates in a "real-time" market (RTM) that attempts to rectify the residual power resultant from forecasting errors and unexpected transients such as generator outages. The prices in the RTM can range from \$30 in one 5-minute window to \$300+ in the next [90], creating significant arbitrage opportunity and fueling energy storage device growth [91]. RTP systems stabilize these fluctuations by bringing more generators and consumers into the real-time market, and these needs will grow with increased renewable integration.

### 5.2.2 Demand-Side Management

Demand side management (DSM) [43, 65] entails a large suite of technologies, including RTP and communication protocols, that bring real-time flexibility to loads in the smart grid. The components in DSMs include digitally controllable loads, network protocols, metering devices, and other pieces required to adapt power supply and demand to a control signal. A general dynamic pricing objective function [73],

used by RTP, is shown in Equation 5.1. The objective is to minimize RP by controlling the market price,  $\lambda$ :

$$\arg \min_{\lambda} \left| \left( \sum_{i \in N} P_i(\lambda, t) \right) \right| \quad (5.1)$$

where  $N$  covers all the consumers/generators in the power grid, and  $P_i(\lambda, t)$  is the power used or produced (negative) by each client at price  $\lambda$  for time  $t$ . Large systems may have multiple  $\lambda$ 's for different locations in the power grid, but this chapter is focused on a smaller market region with a single price signal.

### 5.2.3 Real-Time Communication

Real time communication provides the ability to incorporate dynamic pricing information at the consumer. Distributed consumers have access to changing information that is beneficially incorporated into the pricing optimization problem, i.e.

$$P(\lambda, t) = P_{\text{forecasted}}(t) + P_{\text{flex}}(\lambda, t) + P_{\text{unpredictable}}(t) \quad (5.2)$$

This inherently requires constant communication between the RTP market players to adapt to changes as the system evolves in time. As RTP methods solve for new prices, those prices are broadcast system-wide to the market players. The market players respond by adjusting consumption and production values (i.e. Equation (5.2)) that the RTP algorithm samples, in the physical domain, for its next price calculation.

### RTP Controllers

The core control function samples Equation (5.2) and then solves Equation (5.1). Any number of solutions, from feedback controllers to gradient descent methods, can be used to solve for  $\lambda$ . The information flow in such algorithms may become irregular with imperfect networks, however. In such cases, the market may not perform as expected.

Research in [57] showed that with gradient descent based RTP systems, lost communication messages (e.g. via DoS) would not cause system instability. They assumed that the functions  $P_i(\lambda)$  are stationary and convex—neither of which can be assumed for a distributed system in which the clients can control their own  $P_i$  function. In contrast, work in [78] showed instability for a feedback-based RTP calculations with delays. RTP systems in [3,41,53] rely on (mixed integer) linear programming, gradient descent, interior point, or other optimization techniques to solve for  $\lambda$ . Since these techniques do not account for real-time communication during negotiation phases, technique is utilized and described in Section 5.5.1 to facilitate delay-tolerant solutions to Equation (5.1).

### Impact of Network Outages

Network outages disrupt the communication of  $\lambda$  from the RTP controller to the end users. The  $P_i(\lambda, t)$  for each disrupted  $i$  becomes fixed with  $\lambda = \text{constant}$ , frozen in a zero-order hold. This causes the gain of future  $\lambda$ 's, i.e.  $\frac{\delta P}{\delta \lambda}$ , to decrease since less devices can respond to the change in market price. As a result,  $\lambda$  must go higher or lower to correct for the same amount of RP than in the perfect communication case. To maintain connection incentives, market players are charged retroactively for their power consumption based on the actual market price. This ensures that consumers do not self-disconnect when market conditions appear poor. The residual power measurement is assumed reliable and sampled out of band, e.g. via a dedicated state estimation system.

The strategic adversary is discussed next.

Table 5.1  
List of Symbols

$J$	Set of targets
$\lambda$	Market clearing price (\$)
$T_w$	Attack decision window (s)
$S(t)$	Energy strategy $\in [P_{\min}, P_{\max}]$
$P_i(\lambda)$	Power removed from grid (kW)
$A_j(t)$	Attack target $j$ at time $t$ , $\in 0, 1$
$P_{\text{atk}}$	Estimated price impact of attack (\$)
$D_j$	Flexible load coefficient or gain for target $j$ (kW/\$)
$C_j(t)$	Cost to attack target $j$ at time $t$ (\$)

### 5.3 Strategic Adversary

#### 5.3.1 Strategy Summary

The adversary owns a energy storage device (e.g. a rechargeable battery) and profits by purchasing power at a low price and selling it back into the market at a higher price. To maximize profit, first the adversary monitors the real-time pricing signal and establishes charge and discharge price thresholds. Then she attempts to increase profit by estimating the price impact of a DoS attack by monitoring the gradient of the market price history. Peaks in attack impact are identified as they evolve in real-time, and attacks are launched whenever the peak estimated market price exceeds the charge or discharge thresholds. In this way, the opportunity for arbitrage is maximized, and the adversary increases profitability.

#### 5.3.2 Capabilities and Resources

In this chapter, a model is used where the strategic adversary as a single end user in an RTP system. The adversary can view the price signal and launch attacks against other users in the system.

## User Discovery

Since practical RTP systems are still under development, the attacker is assumed to know the IP addresses of clients participating in RTP market. It is hypothesized, however, that the adversary could discover these addresses in three ways. First, since the market operates on a local level and the potential addresses of Internet-facing devices are geographically correlated for ease of routing, reviewing public IP address registries could reveal targets, especially if RTP participation is widespread. Second, many microgrid applications could utilize peer-to-peer services, especially for islanded operation. These applications could require peer advertisements or open ports that would reveal addresses and service locations. Third, many last-mile network connections utilize shared infrastructure such as cable modem services or passive optical networks. Promiscuous modems could reveal periodic access patterns that are unique to RTP devices, for example. Other alternatives include hacking the RTP controller or other man-in-the-middle security breaches.

## DDoS Attack Capability

The adversary has access to a DDoS-as-a-Service providers or "booter/stressers". Such services offer chunks of attack time for a nominal monthly fee. Armed with a target IP address, the adversary can simply pass it via a web interface and start an attack. It is assumed that consumer grade connections are of sufficiently low capacity such that multiple users can easily be taken offline simultaneously. Additionally, since the attacks are deep in distributed last-mile networks, filtering costs may be prohibitively high. Other smart grid vulnerabilities listed by NESCOR [87] include easy to jam wireless communication channels and physical or logical access to communication channels for entities that do not require it.

## Energy Storage

The adversary is armed with a rechargeable battery that can charge and discharge at a particular rate, has a limited useful lifetime, and a maximum capacity. The battery is assumed to be 100% efficient such that no energy is lost in the charge and discharge process.

### 5.3.3 Strategy Definition

The adversary's arbitrage strategy is to charge the battery when energy is inexpensive and discharge when the price becomes higher:

$$\text{Maximum Revenue} = \arg \max_{S(t)} \sum_{t \in T} T_w \lambda(t) S(t) \quad (5.3)$$

where  $T$  is the set of market clearing windows of negotiation,  $\lambda(t)$  is the market price at time  $t$ , and  $T_w$  is the market clearing or attack strategy window width chosen to discretize the strategy space.  $S(t)$  is the adversary's energy strategy—charge or discharge at time  $t$ . The adversary's goal is to maximize profit by manipulating  $\lambda(t)$  via DoS attacks.

### 5.3.4 Price Manipulation

The adversary can strategically manipulate the market price by launching denial of service (DoS) attacks in the following way. Section 5.2.3 described that whenever clients are disconnected from the marketplace, e.g. via DoS, the effective gain of the price signal decreases. For example, if  $P_i(\lambda) = C\lambda$  for some constant  $C$ , and 10 clients are connected,  $P(\lambda) = 10C\lambda$ . If a DoS attack removes 5 clients, then  $P(\lambda) = 5C\lambda$ , for an attack impact gain of  $5C$ . To achieve the same  $\Delta P$ ,  $\lambda$  would need to change twice as much during the attack. This means that a high RP coupled with client outages leads to more dramatic price swings, and the adversary can leverage these swings to increase revenue. The price manipulation strategy can be broken into estimating RP

and the change in the gain due to a DoS attack to calculate the manipulative power of an attack.

### Target Gain Estimation

In order to influence  $\lambda$ , the adversary needs to know the price gradients of each target's  $P_i(\lambda)$  function. This function is private for each user and unknown even by the RTP service, so the adversary must estimate the gain  $\frac{\delta P_i}{\delta \lambda}$  for each user. One method is to compromise the Internet-connected devices in users home, such as by default passwords or weak encryption, and directly reveal the functions to the adversary. Alternatively the adversary can estimate the gain as  $D_j$  for target  $j$  in the following way. First the adversary measures the gradient of price as a moving average over  $k$  timesteps, e.g. as  $\frac{1}{10C}$ . Then the adversary attacks target  $j$  and measures the new gradient over an additional  $k$  timesteps, e.g. as  $\frac{1}{9C}$ . The gain is then calculated as  $D_j = 10C - 9C = 1C$ . The goal is to observe increases in market prices (or similarly, decreases) and if the attack causes a market participant to go offline, then the rate that the price changes will increase i.e. become convex temporarily ( $|\lambda'_{\text{pre-atk}}| < |\lambda'_{\text{post-atk}}|$ ). This type of approach is not perfect—it is quite possible that targets will be misclassified due to external market conditions such that the price may be concave even without an attack. This classification error simply erodes the adversary's ability to efficiently utilize attack resources (a parameter in experimentation).

### Residual Power Estimation

In the RTP system, the residual power is only known by the RTP controller, as the output of Equation (5.2) and the input of Equation (5.1). The RTP controller decreases and increases  $\lambda$  as a function of RP. Therefore the gradient of  $\lambda$  is loosely proportional to the amount of RP in the system. If RP is negative (shortage), then  $\lambda'$  will be positive, and vice versa for a positive RP. If the gains for every client are

known, then the system gain  $D = \sum D_j$  can be used to estimate RP as  $D\lambda'$ . The attack power is then estimated as

$$P_{\text{atk}}(t) = \lambda'(t) \sum_{j \in J} D_j \quad (5.4)$$

where  $J$  is the set of valid targets,  $\lambda'$  is the current gradient smoothed over  $k$  timesteps, and  $D_j$  is the estimate for  $P'_j(\lambda)$ . It is assumed  $P_{\text{atk}}(t) = 0 \forall S(t) \neq 0$  since the attack is already ongoing. If a net imbalance of power exists, and the RTP signal is actively correcting this by increasing prices, for example, then the attack will cause the price to overshoot by approximately  $P_{\text{atk}}$ .

### 5.3.5 Attack Strategy

Once a set of viable targets and their gains are known, then the adversary may use them to influence  $\lambda(t)$  in an attempt to improve Equation (5.3). First, the baseline strategy is developed with  $\bar{\lambda}_{\text{buy}}$  as the target price for charging periods and  $\bar{\lambda}_{\text{sell}}$  as the target price for discharge. This price is established by a-priori observations of RTP price trends.

Algorithm 2 contains the strategy for attacking targets. Lines 1-3 represent the buying strategy and lines 5-8 the selling one. Line 1 states that whenever the current market price plus the power of attack (which can be negative) is less than the buying price threshold, then the attack should be launched and the battery should charge. Similarly, line 5 sells when the estimated price after attack is higher than the selling threshold. Lines 9-12 stop the attack if the price is not within the buy or sell thresholds. Additional constraints (not shown) keep the battery's energy within capacity. The net benefit from the attack is measured by comparing (5.3) with and without attack for the same scenario.

Algorithm 2 can be further enhanced by peak detection on Lines 1 and 5 rather than operating on the first point that meets the attack standards. The authors of [92]



---

**Algorithm 2:** Basic DoS Strategy
 

---

```

1 if  $\lambda(t) + P_{atk} < \bar{\lambda}_{buy}$  then
2    $A_{j \in J}(t) \leftarrow 1$ 
3    $S(t) \leftarrow P_{\max}$ 
4 end
5 else if  $\lambda(t) + P_{atk} > \bar{\lambda}_{sell}$  then
6    $A_{j \in J}(t) \leftarrow 1$ 
7    $S(t) \leftarrow -P_{\max}$ 
8 end
9 else
10   $A_{j \in J}(t) \leftarrow 0$ 
11   $S(t) \leftarrow 0$ 
12 end

```

---

map real-time peak detection to best choice and optimal stopping problem, and the algorithm is further enhanced by selecting  $\bar{\lambda}$  based upon outlier detection on the  $\lambda(t) + P_{atk}$  signal, as described in [92] Section 3.3.

### 5.3.6 Integrity Attacks

While the focus of this chapter is on network-based attacks, it is also possible to launch integrity attacks on the market. If  $\lambda$  values sent to some subset of clients can be manipulated, then a new attack strategy can be implemented to further defraud the market:

$$P_{atk+}(t) = \sum_{j \in J} \lambda'(t)(P_j(\lambda) - \overline{P_j}) \quad (5.5)$$

$$P_{atk-}(t) = \sum_{j \in J} \lambda'(t)(P_j(\lambda) - \underline{P_j}) \quad (5.6)$$

where  $\overline{P_j}(\lambda), \underline{P_j}(\lambda)$  represent the maximum and minimum power output for each target  $j$ . Algorithm 2 is supplemented by these strategies where Line 1 gets  $P_{atk-}(t)$  and Line 5 gets  $P_{atk+}(t)$ .

Using this strategy, whenever the market price is decreasing due to positive residual power, even more positive residual power is added by further reducing the load

(atk-), causing the market price to plummet further. Similarly, whenever the market price is increasing due to negative residual power, even more load is placed on the system thus increasing price. The net effect is that compromised devices make poor market decisions that benefit the adversary.

### Cost of Attack

The cost of attack can be incorporated into a modified version of (5.3):

$$\arg \max_{S(t)} \sum_{t \in T} \left( T_w \lambda(t) S(t) - \sum_{j \in J} T_w A_j(t) C_j(t) \right) \quad (5.7)$$

where  $A_j(t)$  is the binary attack indicator and  $C_j(t)$  is the cost per second of attacking target  $j$ . The adversary is still attempting to maximize profits in the left term, but each attack that influences  $\lambda(t)$  also has a cost  $= A_j(t)C_j(t)$  in the right term, which can be constrained by a budget (cost  $\leq$  budget).

The costs can be optimized by sorting targets by their cost-impact factors  $C_j D_j$  and prioritizing target above a threshold  $D_{\text{thresh}}$ :

$$A_j = 0 \quad \forall C_j D_j < D_{\text{thresh}}, j \in J \quad (5.8)$$

where  $D_{\text{thresh}}$  eliminates cost-ineffective targets. This restricts the attack strategy by reducing  $P_{\text{atk}}$  to constrain costs.

## 5.4 Defender Strategies

This section covers defensive strategies that can minimize the attacker's profit.

### 5.4.1 Moving Target Defense

The adversary can be countered in two ways. First, the ability to manipulate RP directly can be removed through device security. This is not the focus of this chapter, and is left to other research. Second, the RTP service could remove the ability to disrupt communication links, however the RTP operator would need to harden hundreds or thousands of links to distributed end homes instead of just the network's edges.

Other defensive maneuvers can still be made, however. Intuitively, some targets are "safe" from attack because (5.8) marks them inefficient. The targets  $A_j$  represent IP addresses that will be DDoS'd by the adversary. The defender can mitigate attacks by shuffling the targets ( $A_x \rightarrow A_y$ ) so that the attacker's efforts to attack the IP of target  $x$  actually disable target  $y$ . High value targets can then be swapped for low value targets so that the attacker's profits are minimized. This can be done by synchronizing dynamic IP assignment operations with regional Internet service providers (ISPs)—the RTP operator requests the ISP swap the addresses of  $x$  and  $y$ . The impact is not negligible, however. A forced IP reassignment will cause temporary client interruption, and the creation and support of infrastructure to perform such reassignments would require at least some engineering support. The attacked client would also effectively pay penalty rates for power, so the RTP operator would need to properly incentivize participation.

The defender can strategize about which targets' IP's to swap. Optimally, the highest-value targets would be swapped for the lowest-value targets, and this is what Algorithm 3 performs. The list of targets is sorted by their estimated impact  $D_j$  and the lowest value  $k$  targets are swapped with the highest value targets. This minimizes the change in  $\lambda$  due to  $A_j$

---

**Algorithm 3:** Defender Moving Target Strategy

---

```

1 Sort  $J$  by  $D_j$  descending
2  $k \leftarrow 0$ 
3 while  $k < |J|/2$  AND  $D_{J(k)} > Threshold$  do
4   | Swap IP of  $J(k)$  with  $J(|J| - k)$ 
5 end

```

---

#### 5.4.2 Detection via Deception

The strategies in Section 5.3.3 are all driven from the end-user's observable incentive signal. A false incentive signal could be sent to suspects in the system in order to trigger false attacks on the system. Correlation can be drawn between false signals and corresponding DoS attacks to identify the adversary. For example, a deceptive, high  $\lambda_x$  value could trigger Line 5 in Algorithm 2, and the RTP operator could send this false signal to a potential adversary and observe  $A_j$  via heartbeat signals. This strategy has a cost in that if the signal is sent to a non-adversary, the market efficiency decreases because load will be added or removed contradictory to the current market price  $\lambda$ :

$$\text{Cost of Deception} = |T_w \lambda (P_i(\lambda, t) - P_i(\lambda_x, t))| \quad (5.9)$$

where  $\lambda_x$  is the false price and  $T_w$  is the duration of the false price signal. If deception occurs for one  $T_w$  then the cost is effectively the change in revenue that the client  $i$  was providing to the market. For example, a client uses 1 kWh of energy during  $T_w$ . When the adversary check is run,  $\lambda_x$  is set to  $2\lambda$  and the client consumes 1/4 kWh. The cost is then  $3/4 \cdot \lambda$ .

#### 5.5 Experimental Setup

This section covers the real time pricing mechanism used for experimentation and the load/supply models for evaluation. The full details are covered in Section 3.2

### 5.5.1 Overview of RTP Controller

The price setting algorithm inside the RTP controller minimizes an implicitly non-stationary objective function, the residual power (from Equation (5.1)). Traditional market solutions in [3, 53] utilize gradient descent and interior point methods to solve Equation (5.1), but these techniques require convex objective functions and gradients. Each iteration of an optimization algorithm takes non-zero time, and this has an influence on the behavior of the objective function that may violate necessary assumptions. A stationary  $\lambda$  with a varying  $t$  will change with the supply and demand of electric power, and during transient events such as faults or surge in demand, the change between  $P(\lambda, t - \epsilon)$  and  $P(\lambda, t + \epsilon)$  can be very large. These existing solutions are not equipped to operate in this environment and instead rely on freezing  $t$  for some negotiation period and independently solve for  $\lambda$ . To overcome this limitation, a modified Nelder-Mead (NM) [71] algorithm is utilized that does not make assumptions about the of the objective function and can adapt to large  $\Delta P()$ .

Three modifications of NM are completed to enable the algorithm to perform online optimizations on non-stationary functions. First, the search space used by NM is modified to prevent simplex collapse so that transients can be detected. This is done by adding noise to the points of the simplex so that it maintains a minimum size. Second, the cached function values  $P(\lambda)$  are updated periodically to reflect the current value, and this is used for relative point ranking. This enables the algorithm to adapt to large  $\Delta P()$ . Finally, the algorithm is modified to perform re-evaluation of the simplex space during its shrink operation. Source code for this algorithm is available publicly<sup>1</sup>, including all code used to generate data for this chapter.

### 5.5.2 Load and Generation Model

In this chapter, supply and demand are modeled as scaled sigmoid functions, as shown in Equations 5.10, 5.11 as  $P(\lambda)$ , the power consumed as a function of price

---

<sup>1</sup>[https://github.com/pcwood21/RTP\\_DoS\\_Simulation](https://github.com/pcwood21/RTP_DoS_Simulation)

Table 5.2  
Parameter distributions used in Experiments, respectively for the Consumer (C) and the Generator (G)

	$P_{\min}$ (kW)	$P_{\max}$ (kW)	$\lambda_{\min}$ (\$)	$\lambda_{\max}$ (\$)
C	$ \mathcal{N}(0, 0.5^2) $	$ \mathcal{N}(3, 1^2) $	0	$ \mathcal{N}(250, 75^2) $
G	$- \mathcal{N}(150, 50^2) $	0	$ \mathcal{N}(30, 5^2) $	$ \mathcal{N}(80, 5^2) $

$\lambda$ . Fig. 5.1 shows two example curves with the corresponding residual power at each pricing point, based on the parameters in Table 5.2, where the generator power level is negative. The sigmoid function was chosen to allow responsiveness to price while also ensuring feasible behavior at the extremes, *i.e.*, a consumer cannot have a load greater than  $P_{\max}$  or lower than  $P_{\min}$  and similar constraints for a supplier.

$$\lambda_s = 6 * \frac{\lambda - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} - 3 \quad (5.10)$$

$$P(\lambda) = \frac{P_{\max} - P_{\min}}{1 + e^{\lambda_s}} + P_{\min} \quad (5.11)$$




Fig. 5.1. The supply (generator) and demand (consumer) are shown for varying market prices. The residual power shown is the imbalance between the generator's output and the consumer's input. Since residual power can disrupt the grid and lead to inefficiencies in the energy use, the goal is to optimize the market price such that the residual power is minimized.

For the experiments conducted in this chapter, unless specified otherwise, the parameters used are listed in Table 5.2. There are 2 generators  $G$ ,  $N_G = 2$ , and 100 consumers  $C$ ,  $N_C = 100$ . The adversary's battery has a capacity of 1200 kWh and a charge/discharge rate of 600 kW/h, and  $T_w$  is 5 minutes. The capacity is selected to supplement one generator in the system for 8 hours (e.g. a solar farm during the night).

### 5.5.3 Real-World Dataset

To analyze the effectiveness of the attack strategy at handling unexpected fluctuations in power, the forecasting error from several days of New York Independent System Operator (NYISO) data is used to generate an error function. The difference in the day-ahead forecast and actual load model for June 19-26, 2015 are used to create this signal [90]. The uncontrollable load signal is scaled so that the highest and lowest values are no more than 50% of the maximum and minimum amount that can be absorbed by the generators and consumers in the system. In training, where necessary, the first 7 days are used while the last day is used as the test in all of the experimental results.

## 5.6 Experimental Results

This section covers the results of the evaluation. First the baseline arbitrage opportunities in the market are identified. Then the economic advantage that an adversary can achieve by blocking the communication to a subset of the consumers, first with a naïve attack and then with an attack that tracks the real time price fluctuations, are considered. Then it is evaluated to what degree the attack can be reduced through the defense mechanism. The experiment is concluded by considering the financial gains or investments in dollar terms for the attack.

### 5.6.1 Experiment I: Baseline Profits

In this experiment, the baseline profits are established for any consumer in the attack-free case. As the experiment day progresses, unpredictable changes in consumption cause the price of power to increase and decrease, as shown by the price in Fig. 5.2. The adversary's rechargeable device participates in this market, attempting to minimize the average buy-price and maximize the sell-price to turn a profit.

The charge and discharge duration is limited to two hours by the capacity and charge rate of the battery described in Section 5.5.2, and in this scenario, the adversary is able to profit \$116.48 for the 24 hour period from buying low and selling high. For this strategy,  $\bar{\lambda}_{\text{buy}} = \$64.48$ ,  $\bar{\lambda}_{\text{sell}} = \$146.62$  were chosen via repeated search optimization on the test day. This level of precision is not attainable in practice because the market is assumed unpredictable, so the revenue here is a maximum value using the attack-free strategy in Algorithm 2. This value is not the true maximum possible, since additional arbitrage opportunity exists between hours 20 and 24, but it shows the maximum effectiveness of the heuristic. A more realistic value of  $\bar{\lambda}_{\text{buy}} = \$65.18$ ,  $\bar{\lambda}_{\text{sell}} = \$111.20$ , selected by the top 15% and bottom 15% quantile of prices over the training period, yields a reduced profit of \$70.88. Note that the attack-free strategy is not a harmful event for the power grid—this stabilizes market price and grid loading which is beneficial to consumers and grid operators.

### 5.6.2 Experiment II: Impact of DoS Attacks

In this experiment, the adversary is given the ability to disrupt communication with  $|J| = 20$  users connected to the market. These disruptions increase market volatility by forcing the attacked users to enter a holding pattern in energy consumption. To maximize profits, the adversary implements Algorithm 2. First the adversary selects the parameters  $\bar{\lambda}_{\text{buy}}$ ,  $\bar{\lambda}_{\text{sell}}$ , and she does this by observing  $\lambda(t) + P_{\text{atk}}$  during the training phase. It is assumed that  $D_j$  is known by the adversary, and  $\bar{\lambda}_{\text{buy}} = \$65.00$ ,  $\bar{\lambda}_{\text{sell}} = \$111.35$  are selected by quantiles of 15% and 85% respectively.






Fig. 5.2. The market price during optimal baseline operation is shown. The charge and discharge markers indicate the adversary's optimal charging strategy.

The values for  $D_j, P_{\text{atk}}$  during the test phase are shown in Fig. 5.3. The value of  $D_j$ , shown as the sum over all  $J$  targets, peaks when the market price is relatively low and is suppressed during peak prices. This is because the two largest market players, the generators, are producing maximal output after about \$80. Fig. 3.9 shows that after the \$100 price range, there is only a gradual change in consumption with price increases, mainly by small consumers. In this experiment, the cost of each target is assumed equal, and therefore the power of attack is dominated by the generators.

The adversary implements Algorithm 2 and launches attacks during market operations. Fig. 5.4 shows how the market responds during the attacks. At around 5 hours in, the adversary begins to launch her attacks, and the market price begins dropping in response to these attacks. Once the battery has charged, the attacks end and the market begins to behave normally. After the price rises, the adversary again attacks to increase the market price further. The increase at this price level is smaller due to the low  $D_j$  values in this price range. The attack yields \$119.77 of profit for the day, an increase of 69% over the baseline charging profile.




Fig. 5.3. The power of attack and the gain for 20 targets is shown for the day. The  $D_j$  term becomes saturated at high market prices due to output saturation at the largest market players.




Fig. 5.4. The market price is shown when the attacker implements a DoS attack strategy on 20 targets.

### 5.6.3 Experiment III: Impact of Integrity Attacks

In this experiment, the adversary is given the capability to manipulate individual price signals. The underlying methodology behind the attack is identical to the denial

of service case, but the loads are provided a manipulated price signal instead of a stale one, and the  $P_{\text{atk}}$  is calculated appropriately.

Fig. 5.5 shows impact of compromising devices in the RTP system. The adversary is able to impact price and extract additional profits totalling \$140.36, 98% higher than the baseline. For this attack, the adversary compromises the  $\lambda$  signal as it is sent to the consumer device. The  $P_{\text{atk}}$ ,  $\overline{P_j}$  is achieved by sending  $\lambda = 1000$ , when the adversary wants to buy, and  $\underline{P_j}$  with  $\lambda = -1000$ , when the adversary wants to sell to the grid. This causes the loads imbalance power directly, greatly increasing the effectiveness of the adversary. A cost comparison with DoS attacks is difficult to achieve, however, since the cost of compromising encryption or passwords on consumer devices is not easily quantified.

The effectiveness of these attacks suggests that compromised devices could significantly impede RTP system deployment. Effective defenses, however, are known and need to be deployed more widely, such as, the use of strong authentication scheme and enforcing non-default, and strong passwords. This type of attack can also be very damaging to grid equipment since coordinated loads can cause large transients in voltage and current to occur in the grid, along with instability of RTP systems as has been shown convincingly in [78]. However, that goal is not the focus of this chapter.

#### 5.6.4 Experiment IV: Defensive Strategies

Experiment IV analyzes the defensive strategies presented in Section 5.4. The defender's goal is to reduce or eliminate the adversary's profits. The first defense that the defender implements is protecting the information about the individual consumer loads,  $D_j$ . Investing in stronger end-device encryption and protections, for example, can protect this information. Fig. 5.6 shows how the effectiveness of the attack in Experiment II decreases as the accuracy of the adversary's  $D_j$  terms also decreases. Random noise is added to the  $D_j$  values used in the attacker's strategy to reflect




Fig. 5.5. The market price is shown when the attacker implements a integrity attack strategy on 20 targets. During the charging phase, consumers are misled into conserving power, and during discharge, consumers are misled into over-purchasing power, and this results in price increases and decreases for the adversary to leverage.

inaccurate collection techniques (Section 5.3.4):  $D_j^* = \mathcal{N}(D_j, \sigma^2)$ . The lack of good target information significantly reduces the effectiveness of the adversary, making attacks less profitable. Initially, the adversary's profit drops sharply and then the law of diminishing returns kicks in and the curve flattens out. In this part of the curve, the adversary's estimates are already quite inaccurate and additional noise does not make a significant difference.

Another defensive technique is to swap targets  $D_x$  with  $D_y$ , as described in Section 5.4.1 where the targets are rearranged using Algorithm 3. Fig. 5.7 shows the profit of the adversary versus the number of swaps that the defender is allowed. As the number of swaps increases, the profit decreases but at a lesser rate—the adversary routinely targets the most valuable assets, and since these are first swapped with the least valuable, the impact of the swaps drops off rapidly. When 8 targets are swapped, the adversary's profit is reduced by 30%.




Fig. 5.6. The reduction of the adversary's attack-induced profit is shown as her information about the targets decreases. Errors in target value effectively reduce the profit of the adversary.

### 5.6.5 Return on Investment

The energy storage device used by the adversary, as described in Section 5.3.2 has a particular cost to install and maintain, and the attacks have a particular cost, based on "DDoS/Booter" service pricing [93]. These absolute values can factor in to Equation 5.7 to determine the economic viability of the attack strategies. For revenue, if the adversary repeatedly executes the strategy in Experiment II, amounts could be as much as \$43,000 per year. For cost, storage device prices are expected to fall to \$200 per kWh by 2020 [91] and continue to fall with increases in production, so a yearly battery cost of \$24,000 is estimated and amortized over a 10 year lifetime. The net profit is then  $\$43,000 - \$24,000 = \$19,000$  per year. This cost analysis does not include residual value or added benefits of a distributed battery system such as improved reliability during grid failures. An important note is that if the strategy is profitable, then more devices can yield more profit, or groups of attackers could form battery-consortiums for example. Booter service costs can vary, but average




Fig. 5.7. The reduction of the adversary’s attack-induced profit is shown as targets are swapped by the defender. Since the swaps are optimized on target impact, there is a diminishing return on investment for swapping all of the assets. The first swap protects a large generator with the most impact.

residential Internet connections are low-bandwidth and easily disabled with attacks relative to high-visibility targets like news websites, which keeps these costs low. Based upon the results in [93], a few hundred dollars can maintain a botnet for launching these attacks. For example, 212-booter launched 1993 attacks over 57 days with profits of \$509, suggesting 1,000 attacks could be purchased for about \$250 per month. The end result is that the adversary in this experiment could come out \$15,000 ahead each year with a 1.2 MWh battery.

## 5.7 Discussion

### 5.7.1 Grid Dynamics in Real-Time

Traditional power markets have mostly operated well in advance of actual operation via contracts and optimal schedules. For example, security-constrained unit

commitments (SCUC) [94] have been designed to plan resources with constraints on grid operation contingencies (*e.g.*, to provide stability during generator faults). These systems are still reliant on forecast models, and energy sources that rely on solar and wind will always be subject to some level of uncertainty due to weather patterns. For example, partly cloudy skies can create uncertain solar outputs. This necessitates RTP systems because highly accurate prediction may be implausible. Existing RTP work in [41, 65, 95] has shown the viability of RTP systems to improve efficiency, and a recent study also claimed that dynamic pricing could yield as much as \$400 million per year in savings for the NYISO [64]. These works suggest that the benefits of RTP systems are real, as long as they are properly and securely implemented. This chapter serves to show one way to achieve such secure deployment.

### 5.7.2 Attribution for Attacks

There is some risk that the adversary reveals herself by *a posteriori* profit calculations, *i.e.*, suspicion grows of some market participants that are routinely making gains through price arbitrage or the defender can try to identify the adversary via the technique of injecting false price signals described in Section 5.4.2. Both legitimate consumers and the adversary could respond in similar ways to the market price shifts, however, so plausible deniability may exist from a legal standpoint.

## 5.8 Conclusion

In this chapter, it was presented how a strategic adversary could profit from a real-time pricing system in the smart grid by launching denial of service attacks on consumers connected to the pricing system, *i.e.*, delaying the price signal being sent by the system operator to the consumers. It was shown that an adversary could increase arbitrage revenues by 69% by disrupting up to 20 clients or by as much as 98% if the integrity of the pricing signal is compromised. Then it was shown how a RTP system operator could mitigate network attacks by strategically reconfiguring

the device network. In this way, a defender is able to reduce the adversary's profits from DoS attacks by 30% with 8 IP address swaps. This work exposes some risks to real-time pricing systems in smart grids and provides a novel technique for defending against these attacks.

The next chapter introduces a denial-of-service defense technique that can be used to further reduce the costs of defense for RTP systems.



## 6. DEFENSES

This chapter introduces a denial of service elusion (DoSE) technique for low-cost defense of Internet-facing services. The goal of this chapter is to provide a low-cost framework for protecting RTP operators from attack, especially if the services are regionally operated with a small number of clients. The work presented here also appeared in [96].

### 6.1 Introduction

Denial of Service (DoS) attacks are a continually evolving class of attacks that seek to degrade the ability of legitimate clients to utilize computer resources. Defending against this kind of attack has traditionally been the responsibility of large network operators and internet service providers; however these attacks are increasingly impacting smaller networks or even individual users [97]. The observed trend of increasing attack size, duration, and frequency [98] points to failures in the state-of-practice to mitigate such attacks, especially on an infrastructure level. This is exacerbated by the decreased cost of launching DoS attacks combined with the high relative cost of defending against them with commercial solutions. Time-shared DDoS attacks can be purchased for as little as \$12.99 per month [99] while defense can cost \$2,000 per month [100–102] or more. As long as economic factors favor DoS attacks, they will become increasingly common and persistent occurrences. A low cost solution is needed to prevent stifling of free speech, on the individual side, and to increase the efficiency of doing business, for small to mid-sized businesses.

In this chapter, the beginning effort is presented at achieving the above goal, a system called Denial of Service Elusion (DoSE). Due to the requirement for low economic cost, the defense solution is limited to methods that do not require any

enhancements to the core network infrastructure. Instead, DOSE focuses on using hosting services with widespread and relatively low cost availability such as cloud computing infrastructure and content delivery networks as cornerstones to mitigate attacks. These services have become relatively inexpensive and offer “pay-as-you-go” options which allow flexibility in the mitigation technique. DOSE leverages low cost of public infrastructure-as-a-service (IaaS) cloud and content delivery networks (CDN) to meet an economical cost of roughly \$30 a month for DoS protection for 1,000 clients.

The general approach in DOSE is to connect clients to relays, in an overlay network, instead of directly to a protected service so that DoS attacks cannot easily be launched directly at the service. This technique alone is not novel. However, DOSE adds in a smart management layer which acts to conceal relays from attackers and provide an attack-resistant connection establishment mechanism while most importantly minimizing costs. The relays are created on cloud infrastructure as virtual machines, so the number of relays can expand and contract easily to adapt to changing network conditions. Client-to-relay assignments are communicated over a push-based CDN system that allows for fast reassignments during attack periods, unlike traditional domain name systems (DNS), as well as enabling client-specific assignments. The clients are partitioned among relays and each new relay’s address is selectively released to clients so that if an attack occurs, a set of suspect clients can be identified. The suspects are then separable from the legitimate clients, and with each attack, the suspicion set can be narrowed down. Suspicious clients can be connected to the same relay so that future attacks impact only a small subset of the users.

Prior work in this area is capable of stopping attacks but fails to address the economic considerations of DoS defense. Techniques, such as Portcullis [103] and Epiphany [104], require Internet-wide infrastructure upgrades to combat attacks. Portcullis relies on the assumption that attackers and legitimate clients have similarly-balanced computing power, which may not hold. Epiphany relies on router upgrades (to support reverse multicast) and the availability of thousands of proxy nodes to

defeat DoS (according to *their* experimental setup - Section V-A). MOTAG [105,106] and other overlay network type techniques [107–109] fail to address the resiliency of the client assignment or initial connection channel, relying instead on Portcullis [103] and other existing techniques to stop attacks on a management channel. The work in MOTAG [105] and its subsequent work [106] operate moving target defenses using similar techniques to DOSE. These solutions are not cost-conscious and call for 1,000 active relays for example. Without cost consideration, optimization, or evaluation, they are susceptible to economics-based attacks, whereby the attack exhausts the budget of the consumer for supporting network traffic.

In terms of contributions, DOSE shows how to achieve low cost DDoS attack mitigation for small hosting clients or medium-sized organizations, which have a limited security budget that precludes them from getting a dedicated filtering network or some special arrangements from an ISP. DOSE designs a novel approach for connecting clients to relay proxies and new methods for assigning clients to relays to mitigate network layer attacks while minimizing costs. This work *does not address* attacks capable of disabling large data centers or other large infrastructure networks by well-resourced attackers, or application-layer attacks, i.e., attacks that exhaust the application’s capacity by sending a large number of legitimate-looking, but computationally-expensive-to-process requests.

The rest of the chapter is structured as follows. In Section 6.2, background information is provided on the different kinds of DoS attacks. In Section 6.3, a high-level view is given of the workings of DOSE, followed by the detailed design in Section 6.4. In Section 6.7, the economic costs of using DOSE are laid out versus two existing approaches, one from the commercial domain and the other from the research literature. In Sections 6.5,6.6, the experiments are described with varying numbers and capabilities of legitimate and adversarial nodes and measure what fraction of the non-attack traffic can be supported when the service is under attack. In Section 6.7, the costs of DOSE are analyzed and Section 6.8 presents the conclusions.

## 6.2 Background and Related Work

### 6.2.1 Traditional Defenses

Traditional DDoS defenses technologies rely on filtering or rate limiting traffic along different points of the network [110], such as with access control lists and firewalls. Each technique addresses a way to distinguish legitimate from malicious traffic and then provides a mechanism for increasing goodput [111] by filtering.

The foundation of filtering techniques have a major pitfall, however, because the victim has no control over routers on other autonomous systems (AS) and must rely on cooperation or ingress capacity to begin filtering. To counter this limitation, overlay networks, where traffic is routed to an intermediate server which filters before forwarding on to the destination [107–109], allow defenders to distribute filtering capacity in public clouds without relying on Internet infrastructure support. If the overlay servers are well distributed then a large filtering capacity can be established without modifying any Internet topology devices like firewalls or routers. Another capacity handing mechanism uses anycast [104] that forces traffic destined for a particular victim to be routed to several different servers or networks, each containing the same IP address, so that capacity is improved via redundancy.

### 6.2.2 Overlays and Moving Target Defenses

DOSE builds on two prior techniques, overlay networks and moving target defenses, to achieve low costs. Overlay networks [107–109] provide a layer of indirection to shield protected IP addresses from attacks, and moving target defense techniques control how and where these intermediate computers or relays process traffic to best curtail an attack. With modern cloud computing infrastructure, overlay networks can become elastic and nimble to cheaply dodge attacks, as shown in MOTAG [105, 106] and utilized by DOSE. This elasticity has to be carefully managed, however, to keep the cost of defense low. Simulated attacks in [106] use 1,000 relays over 60 "shuffles"

which each require additional virtual machines to implement. In Amazon’s EC2 service this would translate to as much as 60,000 billable hours of usage, or over \$700 in costs to repel a single attack. The existing moving target methods are stateless and do not account for clients leaving and joining, or the fact that non-aggressive attackers may take hours between subsequent attacks. DoSE focuses on different assignment techniques that minimize these costs by dynamically managing the expenses paid and the number of active relays for any particular defense situation.

## 6.3 DoSE Overview

### 6.3.1 Threat Model

DoSE is targeted at protecting small to medium-sized services, so the adversary is generally comprised of critics and competitors rather than highly skilled cyberwarefare attackers or those with extensive resources. DoSE is designed to defend against attackers that are capable of disrupting communications to a few public IP addresses. DoSE then makes maneuvers to thwart the attacker by making the target unclear and difficult to determine if an attack was successful.

### 6.3.2 Workflow of DoSE

In DoSE, a client does not directly connect to the protected service; in fact, the location of the protected service is kept hidden. Instead the client connects to a relay node on a public cloud infrastructure, which in turn reaches the protected service as explained in Section 6.4.1. DoSE then utilizes a commercial Content Delivery Network (CDN) based system to disseminate the *client-specific* relay information to each client. The dissemination of relay information to a client is done securely through an Assignment Service, with the channel between the Assignment Service and the client being made secure through a shared key that is established as a part of DoSE.




Fig. 6.1. Overview of Infrastructure deployed as part of Dose: The attackers and clients exist on the Internet and are attempting to access a protected service that is either located in a data center or on a small business network (SBN).

A malicious client cannot enumerate the list of relays, preparatory to attacking them, or connect to a relay that it has not been assigned to, as presented in Section 6.4.2.

Figure 6.1 details how the different elements of Dose are laid out. The protected service is in a network which is potentially limited in bandwidth, while the relays are in the network of a public cloud provider (denoted as “datacenter” in the figure). The assignment service, responsible for assigning clients to relays, can be located in either the data center network or the small business network (SBN). A plausible deployment will have all the entities - the protected service, the relays, and the Dose management service – in the cloud environment.

### 6.3.3 Client-Relay Assignment

Whenever a relay is attacked, it can be taken offline and a new relay brought up in its place as in an elastic cloud. This allows the defender to have a seemingly endless supply of fresh relays, whose addresses are selectively released to clients. Further, Dose keeps track of the assignment of clients to the relays so that if a

relay gets overwhelmed due to traffic, a set of suspect clients can be identified. By progressively partitioning suspect clients among multiple relays, DOSE can identify persistent malicious clients as described in Section 6.4.3.

#### **6.3.4 Defense Against Some Obvious Adversaries**

Here some obvious ways are laid out in which an adversary can launch a DoS against the protected service or one of the entities that DOSE introduces.

An attacker identifies the IP address to which she is connected (a relay) and launches an attack at that IP, disabling it and any other clients also connected to that address. Consequently, DOSE creates new relay(s) and re-assigns clients among these new relays. The threat has several characteristics that inform DOSE’s assignment strategy. Each time the client attacks, she is deemed by DOSE to have a higher risk due to prior behavior, i.e., her client being connected to the attacked node, and is therefore managed in such a way to minimize the impact of future attacks as explained in Section 6.4.4.

In the previous attack, the adversary can spawn a new client and try to access the protected service again, concealing her identity. This is handled by the design that a new client is considered inherently “risky”. DOSE will then effectively place new clients in a kind of holding tank to establish trust. Therefore, even under intense attacks, well established and well behaved clients can maintain communications with the service. In the event that an attacker is not aggressive, client connections can be consolidated to a few relays to save costs during non-attack periods and redistributed according to risk when an attack resumes.

#### **6.3.5 Cost Minimization**

Relays are only needed during attack periods. As more clients are involved in attacks, more relays become necessary to maintain communications. The number of

active relays in DoSE is a function of the total connected client risk, controlled by cost parameters via a feedback controller.

## **6.4 Detailed Design of DoSE**

### **6.4.1 Relays**

The relays are simple software filter and forward firewalls on virtual machines in the cloud, mediating client-service communications. Each relay has a set of whitelist firewall rules and rate limiting for each assigned client, and the only entities in the white list are those clients that have been assigned to the relay. For each whitelisted client, there is a rate limit built in at the relay. This can be accomplished with iptables for example.

Attack traffic can come from two sources. The first source is an established client itself, in which case the attacker is known and the excess traffic is stopped by the rate limiting portion of the firewall. The second source is from a non-whitelisted address in which case it is simply dropped. In either case, however, the inbound bandwidth to the relay may become overloaded resulting in a disruption to all of the clients assigned to a particular relay. When an overload occurs, the relay is abandoned and a new relay is established in the cloud. How an attack against a relay node is detected is not central to DoSE, and any monitor which identifies when a relay is inaccessible to the clients is sufficient.

### **6.4.2 Client-Relay Assignment Infrastructure**

The client-relay assignment infrastructure is responsible for communicating the assignment of a client to a relay to the client. It must provide an attack-resilient announcement of IP addresses to those clients while maintaining two properties. First, it must provide each assignment secretly—only the intended client must be able to



read an assignment. Second, it must be able to provide this assignment information to anonymous clients.

To create attack-resilient relay announcements, the assignment is stored on a content distribution network (CDN). These networks store files in a distributed set of caching locations and provide only simple file transfer services (e.g. hash table lookups) which makes them difficult to target at an application layer for attacks. The content itself is replicated across several independent data centers with very fast network connections that make direct infrastructure attacks difficult. The CDN can operate as a push-only replica where the content must be pushed by the provider to the CDN, thus shielding the back end assignment infrastructure. This allows DOSE to withstand large attacks from unauthenticated clients.

This protocol seeks to provide each client with a secure way to acquire its relay assignment using push-only files. At a high level, this is done by having each client solve a puzzle and the solution to the puzzle gives the name of a unique file in the CDN. Following the content in the file, the client is able to contact the relay that it is assigned to.

In the first stage, the CDN has many files with a large number of puzzles and an index, managed by the Identity Establishment Service (IES). A client fetches the index from the CDN and decides randomly to solve one puzzle from the set. The puzzle itself contains an ID, a CAPTCHA image, an integer "PoW Difficulty", and a partial solution hash. The solution to a puzzle comprises of guessing a number from a pre-specified range (PoW) and solving a CAPTCHA, serving as a decryption key.

The client, after solving the CAPTCHA text, guesses a value for the first component of the hash function ("PoW guess"), calculates the hash value, and checks if the hash matches the partial value provided by the puzzle. If it does (say, outcome "A"), it uses the entire hash value as the name of the file to retrieve from the CDN. If it does not (say, outcome "!A"), the client has clearly not guessed the random number correctly. So it tries the next guess. If all guesses are exhausted, then the CAPTCHA was incorrect and a new puzzle must be selected. The scheme forces the client to do

some work before it gains service, limiting request rates. Under outcome A, the client is able to acquire a file from the CDN, call this file “foo1”. This file is itself encrypted using symmetric encryption and the key used is simply  $H(\text{PoW guess}+1, \text{CAPTCHA text})$ . This seemingly contrived design is such that a brute force reading of files from the CDN provides no usable client assignments. The file “foo1” has a client ID and a short-term cryptographic key  $K_s$ . The client ID maps to the name of a file also stored in the CDN, call this file “foo2”. This file has the assignment of the client to a relay node, *i.e.*, reveals the IP address of the relay to the client. During the first interaction with the relay node, made secure using  $K_s$ , the client provides its client ID to the relay node. If, due to collision of the puzzle space, the client ID has already been assigned to a previous client, the relay rejects the initial request from the client. The client then has to go back to the beginning of the process and solve another puzzle en route to getting a new client ID. If, on the other hand, the client ID is unique, then the relay node accepts the request of the client, *i.e.*, forwards it to the ultimate destination. Also, the relay provides a long-term cryptographic key  $K_l$  to the client and to the IES. When the client will be re-assigned to a new relay, say due to an attack, then it will repeat the part of the protocol starting from accessing “foo2” on the CDN. Since it has already acquired a unique client ID, it does not need to repeat the first part of the scheme. For all subsequent interactions with relay nodes, the client will use the key  $K_l$ .

The above scheme for informing a client of its relay assignment achieves the goal that it is using a vanilla CDN, which is “simply” a mechanism to distribute content (files “foo1” and “foo2” in this case). The IES of the DoSE simply pushes content to the CDN and the clients are never allowed to pull information directly from the IES. This provides a level of protection against DoS against the Management service (the IES is part of this) of DoSE.

The “PoW Difficulty” parameter allows the server to control the amount of work it will have the client do before it is allowed to connect to a relay; a higher value implies higher amount of work. The set of puzzles will expire after a certain length of




Fig. 6.2. The client interacts with static files stored on the CDN to retrieve an identity for obtaining future relay assignments in the event of an attack. The client selects a random puzzle from a set on the CDN and then does a proof-of-work to solve the puzzle and retrieve an initial Client ID and relay assignment to begin accessing the protected service.

time, say after  $T$  time units, to prevent solution caching. To reduce the likelihood of collision of the puzzle that a client solves, the number of puzzles in a set will be kept much larger than the number of new clients that are expected to connect to DoSE in time  $T$ . Finally, the client-CDN connection can be established over SSL to prevent man-in-the-middle attacks. An attacker may attempt to solve several puzzles to deny service to new clients by forcing collisions. This attack would require extensive computational resources on the attackers end, especially if the “PoW Difficulty” remains high, resulting in high expenditures for an attacker relative to the IES.

Several conditions must be maintained for this system to be successful. With period  $T$  the content on the CDN is expired and removed. Clients arrive at a rate  $C_r$ , and puzzles are generated at a rate  $P_r$ . Since clients choose puzzles randomly,  $C_r \ll P_r$  ensures a low chance of collision.  $P_r$  is a function of the computation and network capacity of the key management service. A modest capacity of 10 Mbit/s of network upload could provide up to 1000 puzzles per second, assuming a 1 KB image size. The PoW Difficulty parameter allows the server to control the computa-

tional advantage it has over malicious connecting clients at the expense of new client connection establishment times.

### 6.4.3 Client-Relay Assignment Strategy Overview

At the core of DOSE is the many-to-one client-relay assignment strategy in which careful assignment reveals the attacker's identity.

Two categories of adversaries are differentiated. In the first category, the adversary connects, attacks, disconnects, and creates a new identity for herself (as explained in Section 6.4.2). She then launches an identical attack against the newly assigned relay. This type of adversary is a *lone drone*. In the second category, the adversary creates multiple clients concurrently, each assigned to a relay. The clients stay connected to the relay for an extended period, and then, in a coordinated manner, launch a DoS attack against all the relays. This type of adversary is an *aggregate insider*.

In considering the assignment strategy, the notion of *risk* is introduced for each client. Risk is the likelihood that a client will launch a DoS attack in the future. The information items that DOSE currently uses are — length of time the client has used the service (in a well-behaved manner) and whether it is suspected to have been involved in an attack in the past. The latter factor is coarse-grained because when a relay is determined to be under attack, the suspicion falls on *all* the clients assigned to the relay.

New clients are slowly relieved of the *lone drone* risk by the function  $P(\text{lone drone}|t_L) = e^{\frac{-t_L}{EXP(t)}}$  where  $t_L$  is the connection lifetime and  $EXP(t)$  is the expected connection lifetime. If a set of clients is connected to a relay that gets attacked, each client has a uniform probability of being the attacker. The risk assigned to each client is then proportional number of clients connected to a relay. The statefulness of the risk assessment is maintained by recording both the clients involved in each attack and the increase in risk for each client involved in an attack. When an attacker is identified, the risk that was added for all clients implicated in attacks in which the

known attacker was also a suspect, is removed. In summary, the factors that drive the risk estimation of clients is as follows:

- Each client begins with a suspicion of being a lone drone.
- As time progresses, the probability that a client is a lone drone exponentially decays.
- If a client is connected to a relay which comes under attack, its level of suspicion is incremented inverse-proportionally to the number of other clients connected to the relay.
- When an attack is solved, the suspicion of clients associated with the series attacks is reduced.

Note that DOSE still allows for anonymous clients. It only requires there be a secure session between the client and the relay, but it does not require a specific identity of the client to be divulged to the service.

#### 6.4.4 Formal Relay Assignment Strategy

The assignment strategy is designed to minimize percentage of disrupted clients which means evenly spreading risk among each relay.

With the notion of a high risk client being associated with its own relay comes the parameter *risk per relay* (RPR), which is defined as the maximum cumulative risk any relay will tolerate. The total number of relays to use is then calculated as the sum of client risks divided by the RPR. If a set of clients connected to a particular relay is implicated in an attack and this causes the risk on each client to double, say, then the set of connected clients will be split among two relays. Another parameter is the *cumulative risk per attack* (CRPA). This is the total risk increment when a DoS incident occurs. This increment happens for all the clients connected to the attacked relay and the CRPA is divided by that number of clients to arrive at the per-client

increment of risk. The CRPA is parametrized and scaled to control the growth rates of relays. The RPR and CRPA work together to control how the clients are spread during attacks. The higher the  $\frac{CRPA}{RPR}$  ratio, the more quickly the attacker is isolated, but also higher is the number of required relays. If the time to bring a new relay into service increases, then the time to find the malicious client will also increase.

#### 6.4.5 Optimizing Cost

Cost is the direct function of the number of relays utilized and can be controlled by the RPR parameter. A proportional-integral-derivative (PID) controller [112] can be used to adjust this parameter to control long term cost targets at the expense of mitigation effectiveness, specifically by the integral term in the controller.

### 6.5 Experimental Results

Measuring the effectiveness of Denial of Service mitigation is a challenging task, especially when comparing techniques that operate on fundamentally different principles. Prior work [113] suggests using the percentage of failed transactions (*PFT*) to establish effectiveness which is used to measure the success of DoSE.

The effectiveness of different assignment strategies can be assessed without the need for hardware implementations. The actual implementation parameters that influence the performance of DoSE can be easily measured, e.g., relay creation time, iptable filtering rule creation time (at a relay). These performance parameters are then used in an agent-based simulation to evaluate the DoSE approach. While other testbeds exist [114, 115], their goal is to assess filtering methods, which must operate on real-time traffic. In the case of DoSE, once a relay is determined to be unusable due to an attack, it is null routed and henceforth, the actual size of the attack does not have any implication on the network.

### 6.5.1 Management Service Overhead Analysis

The management service has three components that need analysis. The first is the CDN system itself which has been measured sufficiently by prior studies [116, 117]. The second component is the puzzle generation system. Proof-of-work systems are not new, and Portcullis [103] is one example of a network defense implementation which has done performance testing on hash functions. Hashing performance was measured on an EC2 “m1.small” instance to be 1 million hashes per 1.71 seconds, sufficient for DoSE. The next component is the CAPTCHA generation. With 1 ECU, 10,000 GTT’s, 4.2 KB each, were created in 5.11 seconds using libcaptcha. The final component is the assignment and client tracking system. Each session is associated with a key, assignment, and a risk value which is only a small amount of data to store. The assignment algorithm itself is simple — a client is greedily placed on the lowest-risk relay. The conclusion then is that simple EC2 instances can generate a sufficient number of puzzles to supply 10-100 new clients per second with identities.

### 6.5.2 Agent-Based Simulator

To model the DoS scenario, an agent-based simulator [118] is constructed in MATLAB. The simulator facilitates event scheduling and agent-to-agent communications. The agents constructed are the clients, attackers, relay nodes, the assignment service, and the end application. The client is modeled as having a request rate distribution to mimic UDP streaming applications. The attacker has the ability to disable its assigned relay node at any time, causing all traffic to be dropped by that relay.

In the rest of the section, the terms *legitimate clients* and *malicious clients* will be used. Where it can be used without ambiguity, the term *client*, it will refer to a legitimate client. The term *attacker* will be synonymous with malicious client.

### 6.5.3 Experiment 1: Single Adversary

The first experiment has 1,000 legitimate clients and a single attacker to mimic the impact of a DDoS-for-Hire service. Initially all of the clients and the attacker connect at the same time with uniform risk. Then attacker continuously attacks, and the experiment ends once the malicious client is isolated.

Figure 6.3 shows the average PFT for the set of connected clients. In this case, the minimum number of relays is two. After the initial attack, one relay is disabled, resulting in a PFT of 50%. Additional relays are brought online to distribute the now high risk set of disrupted clients. This continues until 3.9 minutes into the simulation when the attacker is identified. The time needed to neutralize the malicious client depends on the cumulative risk per attack (CRPA) parameter and the time to bring a relay online (40 seconds), as described in Section 6.4.4. Over the time span of the entire experiment, the average PFT is 0.118, meaning that an average of 88.2% of the requests are satisfied during the mitigation period.

Another important metric is the number of relays used in the defense. Figure 6.4 details how the number of relays grows with time, growing from an initial 2 relays to 16 relays. The growth rate is a function of the  $\frac{CRPA}{RPR}$  ratio. After each attack cycle, additional relays are created to reduce the impact of attacks and prune the suspect list. The number of relays is also dependent on the number of clients connected. Figure 6.5 shows the number of relays in use during the same attack but with only 100 clients. In this case, the attacker is found more quickly and the overall usage is lower, growing to a maximum of 11 relays. This drives home the point that costs in DoSE are client-dependent. Note however, that there is no strict proportionality between the number of clients and the number of relays needed to mitigate the attack. The formulae used in the algorithms for assignment of clients to relays are non-linear and this explains the observation.






Fig. 6.3. Percentage of Failed Transactions with a single attacker: The attack begins at 80 seconds and is progressively mitigated with each relay-creation cycle until the attacker is identified and neutralized at 3.9 minutes. The relay power-on time is 40 seconds, the approximate width of the steps.

#### 6.5.4 Experiment 2: Streaming Attack

In the second experiment, the stress due to the attack is higher with multiple malicious clients launching coordinated attacks. This is referred to as the *Streaming Attack*, consisting of *aggregate insiders*. New legitimate clients connect to the service uniformly every 1.6 seconds and stay connected for 400 seconds, averaging 250 connected clients. The attackers arrive in sets of 10 every 160 seconds and simultaneously attack. Eventually, the malicious clients are identified and neutralized.

Figure 6.6 shows how the PFT is impacted by the streaming attack. The initial attack is difficult to mitigate because all clients have uniform risk, few clients are legitimate, and the relay count is small. As the client base matures, subsequent attacks have a much lower impact on the PFT because the new clients are high risk compared to existing, legitimate ones. It is seen that even the spikes, which




Fig. 6.4. Number of Relays Used: Defending against a single attacker with 1000 clients shows increased relay counts until the attacker is found.

correspond to the arrival of the new batch of malicious nodes, only go up to 20% failed transactions.

### 6.5.5 Experiment 3: Attackers Present at Startup

In this experiment, the case when the malicious clients are present at initialization is evaluated. The performance of DOSE is evaluated as the number of malicious clients is varied from 1 to 500, with a constant legitimate client base of 1,000. The PFT and relay count metrics are averaged over 800 seconds, by which time the last malicious client has been identified and isolated.

Figure 6.7 shows how many relays are used to defeat attacks of varying sizes of number of malicious clients. With a large number of attackers, the number of relays required to distinguish between attackers and legitimate clients becomes large as well. With 1 malicious client, the number of relays required is 6; with 15, it is 30; and, with




Fig. 6.5. Number of Relays Used: Defending against a single attacker with a smaller number of clients shows the relay count to be smaller than with 1,000 clients.

500 malicious clients, the number of relays required is 154. There are two factors to note here. First, the ratio of adversarial to legitimate clients of 1:2 is rather high, and second, the growth in the number of relays is slower than the increase in attackers. Also, an insider and a legitimate client begin with indistinguishable features and history. A more likely case is that some clients will have a higher trust due to a long association with the service.

Figure 6.8 details the impact on the average PFT during the attacks. Expectedly, as the number of attackers increases, the PFT increases. As each attack occurs, the legitimate clients are prevented from accessing the service until a relay is assigned only legitimate clients. With this level of discrimination, the malicious client(s) can be isolated. However the convergence time to reach this point increases with the number of attackers, thus driving up the average PFT.




Fig. 6.6. Percentage of Failed Transactions with the streaming attack: The attack begins at 160 seconds and is repeated with 10 malicious clients arriving every 160 seconds. A legitimate client arrives every 1.6 seconds and stays connected for 400 seconds. As time progresses, the risk profile of the legitimate clients decreases leading to a better isolation of the malicious nodes. Consequently, the PFT decreases compared to the initial burst of attack.

## 6.6 Result Analysis

### 6.6.1 Sensitivity to CRPA

The CRPA parameter, as described in Section 6.4.4, controls the amount of relay growth per attacked relay. If the CRPA is 2 and the RPR is 1, then a single attacked relay will be replaced by 2 new relays. Figure 6.9 shows the relay growth from Experiment 1 (single attacker, Section 6.5.3), with varying CRPA factors.

In this experiment, the attacker is flooding the relays continuously and therefore, the high CRPA value gives the best result—the time to identify the attacker is the lowest (compared to lower CRPA values) and the total cost is the same. The downside of a high CRPA value will be exposed if the attacker is more subtle, and after




Fig. 6.7. Number of Relays vs Number of Malicious Nodes: The number of legitimate clients is kept fixed at 1,000 and the number of malicious clients is increased. All the clients are present at startup. The increase in the number of malicious clients is met with automatic and progressive increases in the number of relays, to isolate the malicious nodes.

launching an attack (and causing a large number of relays to be created), it lies low for a while. The spurt in the number of relays does not help DOSE if condensed prior to a subsequent attack.

### 6.6.2 Quantitative Comparison with Epiphany and Speak-up

A quantitative comparison of Epiphany [104] and Epiphany is done, coupled with Speak-up [119] with DOSE for the setup of experiment 3. DOSE is executed and the number of relays used is given as the number of proxies in Epiphany. To do a comparison, a distribution was created of clients and attackers on the Internet, such that 10% of the Internet is “unclean networks”, where 80% of the IP addresses are attackers. The remaining 90% of the Internet is split between completely clean and partially clean networks, with the clean network accounting for 10% of the addresses,




Fig. 6.8. PFT vs Number of Malicious Clients: The number of legitimate clients is kept fixed at 1,000 and the number of malicious nodes is increased. All the clients are present at startup. The increase in the number of attackers causes the average PFT over the attack simulation window to grow as a larger-sized attack impacts more client-connected relays.

defined as having no malicious address. The partially clean network has 20% of its addresses as attackers. The design of Epiphany suggests that if a legitimate client is in the unclean network, then she will not be able to have any successful transactions. If the client is assigned to a clean network, then there will be no impacts of the attack. Therefore the percentage of failed transaction varies between 10% and 90%. The remaining clients may or may not be connected to proxies which contain attackers. If there is an attacker, then the ratio of successful transactions is the percentage of good clients connected. Thus if there are 9 clients connected to a proxy and 1 attacker, 90% of the transactions will be successful, modelling the behavior of Speak-Up + Epiphany [119]. In the alternative mode, without Speak-Up, any client assigned to a relay with an attacker will not be able to access the protected service.




Fig. 6.9. Number of relays over time for varying CRPA factors: The number of relays utilized peaks to a much higher number with high CRPA values, however the attacker is found much more quickly. The factor controls the growth rate in the number of relays created in response to an attack.

Figure 6.10 shows the results of the comparison. At a low number of attackers, DoSE is able to maintain the advantage. At higher numbers of attackers, Epiphany + Speak-Up is able maintain the best advantage while DoSE outperforms the pure Epiphany solution. However, as pointed out earlier, DoSE requires no widespread network infrastructure upgrades.

### 6.6.3 Comparison with MOTAG

MOTAG [105] provides an alternative approach to client-relay assignments from DoSE. It uses a greedy assignment algorithm in an attempt to save as many clients as quickly as possible. To make a comparison between DoSE and MOTAG, DoSE was run and the average number of relays used during the scenario was used as the input to MOTAG. MOTAG also relies on a set of shuffling proxies and a set of




Fig. 6.10. Comparison of DoSE to Epiphany and Epiphany+Speakup: At low ratio of number of attackers to legitimate nodes, DoSE outperforms other solutions, while at mid to high ratios, Epiphany+Speak-up outperforms DoSE.

serving proxies, and legitimate clients are moved from the shuffling to the serving proxies as the attack progresses. It is assumed that there is a distribution of half and half between the shuffling and the serving proxies. Additionally, MOTAG provides an estimation technique for the number of insiders; here, the truth value is simply provided to MOTAG. In this evaluation, a rate of one IP address is being revealed per minute so multiple insiders are independent sources of addresses for attack.

Figure 6.11 shows that MOTAG takes much longer to isolate the attackers than does DoSE. This is because MOTAG relies on stateless transitions between attacks while DoSE builds a long-standing history, through the risk value, for each client. Therefore an insider who stops attacking is not forgotten, while in MOTAG, the insiders may continue to impact legitimate clients if moved to a serving proxy. A higher time to find the attacker results in increased costs since a larger number of relays are ultimately used in defense, thus DoSE costs 59% less than MOTAG. Figure 6.12 shows the number of failed transactions per second, averaged across the total






Fig. 6.11. Comparison of DoSE to MOTAG. The time it takes to find all of the insiders is significantly lower in DoSE, especially as the number of insiders increases. This is due to the smart relay management which divides the clients during successive attacks in DoSE while MOTAG relies on a stateless classification which lets intelligent adversaries fool the system.

simulation time. Since DoSE finds the attacker more quickly than does MOTAG, fewer transactions fail when using DoSE to defend against this attack scenario.

## 6.7 Costs and Discussion

The costs of DoSE are application dependent with two components: time-of-use resources (relay instance runtime) and bandwidth-based network resources (cost to forward). Relay nodes on EC2 [120], for example, cost \$8.64 a month for the smallest instance or \$2.23 with spot pricing. Network resources are a function of legitimate client traffic that must be passed from the cloud to the protected service. Costs for forwarding traffic out-of-network are priced per GB and range from \$0.05 to \$0.12 per GB, but can be mostly avoided by co-locating the relays and protected service.




Fig. 6.12. Comparison of DoSE to MOTAG. DoSE is able to keep the number of failed transactions lower during the simulation because it better isolates and identifies the attacker more quickly, resulting in less clients co-located on relays with insiders.

A comparison can be made between the cost of DoSE and existing commercial solutions. For a representative sample, Staminus Communications [100] will be used. The defensive cost is related to the capability of the adversary, since a filtering technique is used, and the cost comes to about \$35 per Gbit, \$175 per MPPS of attack protection. The defensive cost for an average attack of 9.7 Gbit/s, 19.8 MPPS peak [98] would be \$3,465 per month.

The cost of DoSE can be compared to Epiphany [104] and MOTAG [105, 106]. For the former, a set of static proxies are needed as entry points into the Epiphany routing structure. The cost of operating a proxy for Epiphany will be 50% higher than EC2 because it is transitioning from a datacenter style location to a more end-user style location (drawing from Amazon’s differentiated tiers). To estimate the router upgrade cost, if each router takes 1 man-hour to reconfigure, and there are 1 million routers in the system, then upgrading 0.01% (a parameter chosen in [104])

of the routers to support Epiphany would require 1,000 man-hours or \$100,000 at \$100/hour for a network engineer. If there are 10 routers between the source and destination and each router has 10 output ports enabled, the amount of bandwidth used will be 100 times the original transmission due to reverse-multicast.

For MOTAG, cost-equivalent performance comparisons can be seen in Section 6.6.3. As tested in [106], MOTAG relied on 1000 replicas through 60 shuffles which amounts to 60,000 billable hours (at 1 hour minimums) on EC2, or over \$700 to stop a single attack. With that many replicas in play, more efficient static approaches could keep clients connected.

These findings are summarized in the following table using the shorthand “R” for relay, “m” for month, and R-dependent means it depends on the number of relays.

Defense	Overhead	BW	Atk. Size
DoSE	\$9/R/m	2x	Independent
Epiphany	\$22/R+\$1670/m	100x	R-Dependent
Commercial	\$3,465/m	Included	100 Gbps/20 MPPS

Another cost of using any design as in DoSE of separate entities—relay and protected service—is the latency penalty for taking alternative routes in the network. This has been measured in [105], and a round trip time penalty of 70 ms would be typical for the continental US.

The cost of DoSE for the experiments in Section 6.5 can be calculated for EC2. In the first attack example 6.5.3, a maximum of 16 relays are used, and there are 5 rounds of relay expansion captured by the number of steps seen in the graph. Since EC2 requires a minimum billing of 1 hour for each instance started, the 5 instances that were disabled by the attacker along with the 16 relays consumed are billed for 1 hour, totaling 21 hours of billing on the smallest EC2 instance. The cost is then \$0.42 for on-demand instances or \$0.07 for spot instances to stop this attack. Additional costs will depend on how much Internet bound traffic passes through the relay.

## 6.8 Conclusion

DOSE was presented as a method for mitigating network-layer Denial of Service attacks. DOSE uses cloud-based relay nodes to hide protected services from direct attack while acting as sacrificial targets. The key idea behind DOSE lies in its ability to assign relays to specific clients and by re-assigning suspect clients to a progressively smaller set of relays, to identify and isolate the malicious clients. DOSE also smartly manages the set of relays to optimize cost while allowing maximal attack mitigation. The technique is capable of quickly mitigating attacks while continually improving on the legitimate client's ability to complete transactions. It is believed that DOSE is the first technique that can achieve DoS protection for a price point that would be acceptable to medium-to-small-sized businesses, of less than \$100 per month.

## 7. CONCLUSION AND FUTURE WORK

### 7.1 Conclusion

Cyber-physical systems, especially those that control electric power consumption, will continue to expand and grow as new control models such as real-time pricing continue to improve the efficiency of a variety of industrial processes. These systems, however, may be vulnerable to availability attacks via network disruptions. In this dissertation, Chapter 2 introduced the energy-specific usage cases for cyber-physical systems. Renewable energy resources and market deregulation motivated the creation of a multi-agent profit-oriented model for strategic adversaries. Chapter 3 introduced the technical underpinnings of operating such a CPS via real-time pricing systems and dynamic market mechanisms. The impacts of latency were also examined in that chapter. Chapter 4 combined the previous two chapters into a attack/defense game for launching network attacks on an real-time pricing system. The results validated the work in this dissertation, but it introduce some shortcomings in the strategy space due to its large size. Chapter 5 then demonstrated a technique for attacking systems in real-time with a charge/discharge strategy and denial of service attacks. Finally, Chapter 6 provided a denial-of-service defense with low cost optimizations. This technique can be combined with the RTP system to achieve low-cost resilient market operations.

The primary contributions of this dissertation include:

- A multi-agent model for energy-based cyber-physical systems where one or more agents are malicious
- An attack and defense strategy for those CPS executed under MILP

- An analysis of low-level network attacks and their mapping to the attack/defense strategy space, including defensive maneuvers
- An analysis of latency impacts on two CPS control models, including a novel control system
- An on-line attack heuristic for extracting profits from a real-time pricing system
- A low-cost denial-of-service defense technique (DoSE) for lowering the cost of defending such systems

In summary, this dissertation has laid the groundwork for engineering resilient cyber-physical systems from the perspective of a strategic adversary. As CPS grow and include multiple financially independent entities under their control, financial incentives will need to exist to promote cooperation. Such systems, like real-time pricing, may be subject to exploitation by strategic adversaries as shown in this dissertation.

## **7.2 Future Work**

Several future directions for additional work include:

### **7.2.1 Multi-Round Attacker/Defender Games**

The game theoretic models in Chapters 2 and 4 are designed as single round games. The defender first strategies on defense and then the adversary makes a move based on static knowledge about the system. This model can be improved with multi-round games that allow the adversary to adapt to the defenders' moves. Such a system would improve the fidelity of both the attack and defender models and would provide additional insight into asset protection schemes and cost sharing models.

### 7.2.2 Stateful Real-Time Attacks

The attack models covered in Chapters 4 and 5 are based on stateless consumers and producers. This means that their loads and needs are modeled as exogenous functions that are not dependent on what happens in the operational system. Practically, these producers and consumers operate with time-shiftable loads, fixed fuel storage, etc. that would cause stateful responses in the power grid. For example, if a consumer avoids loading the grid at  $t = 10$  h due to market conditions, then that consumer would have a much higher need to load the grid at  $t = 11$  h. The implication is that the fidelity of the post-attack response of the grid can be improved by including stateful client models in the RTP system. This future work would enhance the adversary's profit if properly executed.

### 7.2.3 Client-Side DoSE

DoSE, presented in Chapter 6, is designed to protect a single service with multiple clients from attacks. Some RTP systems, and those that are designed in the future, exist where the clients are the attack target rather than the central service. Such attacks require new overlay network design models to protect clients from attacks, e.g. via intra-overlay routing. Methods that secure this kind of interaction at low cost could protect the systems described in Chapter 3.

### 7.2.4 Algorithm Resilience

This dissertation focuses on mitigating attacks by undoing their effect (i.e. stopping a DoS attack). An alternative or supplemental approach is to design the algorithms to operate in a degraded state such that the impact of attack is minimized. For example, in the market mechanisms presented in this dissertation, the clients enter a zero-order hold state after attack. Instead, clients could adopt some contingency plan or some other operational state that reduces the adversary's profit, etc. Such

work would further improve the defensive strategies that could be used in the game theoretic planning phase of system design.



## APPENDICES

## A. SIMULATION ENVIRONMENT

Throughout this dissertation, several simulation tools were utilized to produce experimental results. This appendix describes those simulation tools and provides links to source code where available. Additional source code not listed in this appendix can be retrieved by contracting the Dependable Computing System Lab<sup>1</sup>.

### A.1 Game Theoretic Model for CPS

Chapter 2 contains several computational models for calculating the optimal attack/defense strategies in CPS. This model is available publicly<sup>2</sup> or via DCSL. The model exists as MATLAB code with several functions designed to create and solve the intermediate steps of generating the impact matrix and using it to drive strategies. The folder `IDD_Model` contains the information about the experiment, and the folder `cluster_scripts` contains several examples on how to generate results from this model.

### A.2 Agent-Based Model

Several models in this dissertation are based on an agent-based framework in MATLAB<sup>3</sup>. This framework contains an event scheduler, agent models, and a mechanism by which agents can schedule events and communicate with each other. A logging mechanism is also included in this model. An additional batch script tool is available<sup>4</sup> that allows jobs to be submitted to a cluster for quick evaluation of various simulation scenarios. The subsections included all utilize this model.

<sup>1</sup><https://engineering.purdue.edu/dcs1/>

<sup>2</sup><https://github.com/pcwood21/CPSSim>

<sup>3</sup><https://github.com/pcwood21/DSim>

<sup>4</sup><https://github.com/pcwood21/BatchScriptDSIM>

### A.2.1 Technical Markets

Chapter 3 covers two market models, DMM and NM. The former is not available publicly because a core component of the source code (the DMM itself) is not yet released. The Nelder-Mead model and its supporting framework is available publicly<sup>5</sup>, including the code to generate figures used in this dissertation. This model utilizes the agent-based framework with a market operator and connected clients.

### A.2.2 Technical Markets and Game Theory

Chapter 4 utilized NM with the work in Chapter 2. The source code for this is available publicly<sup>6</sup>. The source code under the Models folder in this repository can be used to re-create all of the figures in the experiment.

### A.2.3 Real-Time Market Attacks

Chapter 5 utilized the NM market mechanism with an on-line attack heuristic. This code is available<sup>7</sup>, but it does not utilize the agent-based scheduling framework provided. Instead, function calls are made directly to the objects to reduce execution overhead.

### A.2.4 DoSE Model

The experiments in Chapter 6 are based on a publicly available model<sup>8</sup>. The agents include the end clients, the protected service, and the assignment service. The CDN is not modeled directly, and client assignments are assumed to always be transmitted appropriately. Additional comparison models are in the source code as well.

---

<sup>5</sup>[https://github.com/pcwood21/NM\\_Market\\_Model](https://github.com/pcwood21/NM_Market_Model)

<sup>6</sup>[https://github.com/pcwood21/CPS\\_Model\\_Comsnet16](https://github.com/pcwood21/CPS_Model_Comsnet16)

<sup>7</sup>[https://github.com/pcwood21/RTP\\_DoS\\_Simulation](https://github.com/pcwood21/RTP_DoS_Simulation)

<sup>8</sup>[https://github.com/pcwood21/DoSE\\_ABM](https://github.com/pcwood21/DoSE_ABM)

## LIST OF REFERENCES

## LIST OF REFERENCES

- [1] F. C. Schweppe, R. D. Tabors, J. L. Kirtley Jr, H. R. Outhred, F. H. Pickel, and A. J. Cox, "Homeostatic utility control," *Power Apparatus and Systems, IEEE Transactions on*, no. 3, pp. 1151–1163, 1980.
- [2] P. Fairley, "Innovation amid a raucous rooftop solar squabble," *Spectrum, IEEE*, vol. 52, pp. 14–15, July 2015.
- [3] D. Shiltz, M. Cvetkovic, and A. M. Annaswamy, "An integrated dynamic market mechanism for real-time markets and frequency regulation," *IEEE Transactions on Sustainable Energy*, 2015. to appear.
- [4] P. Wood, S. Bagchi, and A. Hussain, "Optimizing defensive investments in energy-based cyber-physical systems," in *Parallel & Distributed Processing Symposium Workshops (IPDPSW), 2015 IEEE International*, IEEE, 2015.
- [5] C.-W. Ten, G. Manimaran, and C.-C. Liu, "Cybersecurity for critical infrastructures: Attack and defense modeling," *Trans. Sys. Man Cyber. Part A*, vol. 40, pp. 853–865, July 2010.
- [6] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *Security & Privacy, IEEE*, vol. 9, no. 3, pp. 49–51, 2011.
- [7] ICS-CERT, "Ics-cert monitor, internet accessible control systems at risk," 2014.
- [8] T. Egan, "Tapes show enron arranged plant shutdown," 2005.
- [9] C. Bronk, "Hacks on gas: Energy, cybersecurity, and u.s. defense," tech. rep., Baker Institute for Public Policy, 2014.
- [10] H. Pinto, F. Magnago, S. Brignone, O. Alsac, and B. Stott, "Security constrained unit commitment: network modeling and solution issues," in *Power Systems Conference and Exposition, 2006. PSCE'06. 2006 IEEE PES*, pp. 1759–1766, IEEE, 2006.
- [11] M. Shahidehpour, Y. Fu, and T. Wiedman, "Impact of natural gas infrastructure on electric power systems," *Proceedings of the IEEE*, vol. 93, no. 5, pp. 1042–1056, 2005.
- [12] M. Urbina and Z. Li, "A combined model for analyzing the interdependency of electrical and gas systems," in *Power Symposium, 2007. NAPS'07. 39th North American*, pp. 468–472, IEEE, 2007.
- [13] C. Correa-Posada and P. Sanchez-Martin, "Security-constrained optimal power and natural-gas flow," *Power Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1–8, 2014.

- [14] T. Li, M. Eremia, and M. Shahidehpour, "Interdependency of natural gas network and power system security," *Power Systems, IEEE Transactions on*, vol. 23, no. 4, pp. 1817–1824, 2008.
- [15] L. Wang, S. Noel, and S. Jajodia, "Minimum-cost network hardening using attack graphs," *Computer Communications*, vol. 29, no. 18, pp. 3812–3824, 2006.
- [16] R. Dewri, N. Poolsappasit, I. Ray, and D. Whitley, "Optimal security hardening using multi-objective optimization on attack tree models of networks," in *Proceedings of the 14th ACM conference on Computer and communications security*, pp. 204–213, ACM, 2007.
- [17] K. Wang, B.-h. Zhang, Z. Zhang, X.-g. Yin, and B. Wang, "An electrical betweenness approach for vulnerability assessment of power grids considering the capacity of generators and load," *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 23, pp. 4692–4701, 2011.
- [18] P. Hines, E. Cotilla-Sanchez, and S. Blumsack, "Do topological models provide good information about electricity infrastructure vulnerability?," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 20, no. 3, p. 033122, 2010.
- [19] T. M. Chen, 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. 4, pp. 741–749, 2011.
- [20] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: an overview of microgrid systems, demand-side management, and smart grid communications," *Signal Processing Magazine, IEEE*, vol. 29, no. 5, pp. 86–105, 2012.
- [21] K. J. Ross, "Application of game theory to improve the defense of the smart grid," tech. rep., DTIC Document, 2012.
- [22] S. Backhaus, R. Bent, J. Bono, R. Lee, B. Tracey, D. Wolpert, D. Xie, and Y. Yildiz, "Cyber-physical security: A game theory model of humans interacting over control systems," *Smart Grid, IEEE Transactions on*, vol. 4, no. 4, pp. 2320–2327, 2013.
- [23] S. Amin, G. A. Schwartz, and A. Hussain, "In quest of benchmarking security risks to cyber-physical systems," *Network, IEEE*, vol. 27, no. 1, pp. 19–24, 2013.
- [24] Y. Mo, T.-H. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, "Cyber-physical security of a smart grid infrastructure," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 195–209, 2012.
- [25] A. Laszka, M. Felegyhazi, and L. Buttyan, "A survey of interdependent information security games," *ACM Comput. Surv.*, vol. 47, pp. 23:1–23:38, Aug. 2014.
- [26] M. H. Manshaei, Q. Zhu, T. Alpcan, T. Başar, and J.-P. Hubaux, "Game theory meets network security and privacy," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, p. 25, 2013.

- [27] H. Chan, M. Ceyko, and L. E. Ortiz, "Interdependent defense games: Modeling interdependent security under deliberate attacks," *arXiv preprint arXiv:1210.4838*, 2012.
- [28] V. M. Bier, "Choosing what to protect," *Risk Analysis*, vol. 27, no. 3, pp. 607–620, 2007.
- [29] S. Amin, G. A. Schwartz, and S. Shankar Sastry, "Security of interdependent and identical networked control systems," *Automatica*, vol. 49, no. 1, pp. 186–192, 2013.
- [30] K. Hausken, "Income, interdependence, and substitution effects affecting incentives for security investment," *Journal of Accounting and Public Policy*, vol. 25, no. 6, pp. 629–665, 2006.
- [31] M. A. Delucchi and M. Z. Jacobson, "Providing all global energy with wind, water, and solar power, part ii: Reliability, system and transmission costs, and policies," *Energy Policy*, vol. 39, no. 3, pp. 1170–1190, 2011.
- [32] U.S. Energy Information Administration, "Natural gas," 2014. <http://www.eia.gov/naturalgas/>.
- [33] U.S. Energy Information Administration, "Electricity," 2014. <http://www.eia.gov/electricity/>.
- [34] U.S. Federal Energy Regulatory Commission, "FERC: Natural gas," 2014. <http://www.ferc.gov/industries/gas/gen-info/fastr/index.asp>.
- [35] J. M. Petrash, "Long-term natural gas contracts: Dead, dying, or merely resting," *Energy LJ*, vol. 27, p. 545, 2006.
- [36] J. Hansen, J. Knudsen, and A. M. Annaswamy, "A dynamic market mechanism for integration of renewables and demand response," *Control Systems Technology, IEEE Transactions on*, 2015. to appear.
- [37] X. Zhang and A. Papachristodoulou, "A real-time control framework for smart power networks: design methodology and stability," *Automatica*, vol. 58, pp. 43–50, 2015.
- [38] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Dynamic network energy management via proximal message passing," *Foundations and Trends in Optimization*, vol. 1, pp. 70–122, 2013.
- [39] X. Zhang, N. Li, and A. Papachristodoulou, "Achieving real-time economic dispatch in power networks via a saddle point design approach," in *IEEE PES General Meeting, Proceedings of*, 2015.
- [40] E. Mallada and S. Low, "Distributed frequency-preserving optimal load control," in *IFAC World Congress, Proceedings of*, pp. 5411–5418, 2014.
- [41] P. Samadi, A.-H. Mohsenian-Rad, R. Schober, V. W. Wong, and J. Jatskevich, "Optimal real-time pricing algorithm based on utility maximization for smart grid," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pp. 415–420, IEEE, 2010.

- [42] X. Fang, S. Misra, G. Xue, and D. Yang, “Smart grid: the new and improved power grid: A survey,” *Communications Surveys & Tutorials, IEEE*, vol. 14, no. 4, pp. 944–980, 2012.
- [43] M. Petersen, K. Edlund, L. H. Hansen, J. Bendtsen, and J. Stoustrup, “A taxonomy for modeling flexibility and a computationally efficient algorithm for dispatch in smart grids,” in *American Control Conference (ACC), Proceedings of*, pp. 1150–1156, IEEE, 2013.
- [44] M. Kuzlu, M. Pipattanasomporn, and S. Rahman, “Communication network requirements for major smart grid applications in HAN, NAN and WAN,” *Computer Networks*, vol. 67, pp. 74–88, 2014.
- [45] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot, “Characterization of failures in an IP backbone,” in *IEEE INFOCOM, Proceedings of*, vol. 4, pp. 2307–2317, Mar. 2004.
- [46] M. Ilic, J.-Y. Joo, P. Carvalho, L. Ferreira, and B. Almeida, “Dynamic monitoring and decision systems (DYMONDS) framework for reliable and efficient congestion management in smart distribution grids,” in *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pp. 1–9, Aug. 2013.
- [47] A. Hahn and M. Govindarasu, “Cyber attack exposure evaluation framework for the smart grid,” *Smart Grid, IEEE Transactions on*, vol. 2, pp. 835–843, Dec. 2011.
- [48] D. Shelar and S. Amin, “Analyzing vulnerability of electricity distribution networks to DER disruptions,” in *American Control Conference (ACC), Proceedings of*, pp. 2461–2468, July 2015.
- [49] N. A. E. R. C. (NERC), “Reliability concepts technical report v1.0.2,” tech. rep., North American Electric Reliability Corporation (NERC), Dec. 2007.
- [50] D. Callaway and I. Hiskins, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, pp. 184–199, Jan. 2011.
- [51] A. M. Annaswamy and T. R. Nudell, “Transactive control—whats in a name?,” *IEEE Smart Grid Newsletter*, Sep. 2015. <http://smartgrid.ieee.org/newsletter/september-2015/transactive-control-what-s-in-a-name>.
- [52] M. Ilic and J. Zaborszky, *Dynamics and Control of Large Electric Power Systems*. New York: John Wiley and Sons, 2000.
- [53] R. Zhou, Z. Li, and C. Wu, “An online procurement auction for power demand response in storage-assisted smart grids,” in *IEEE INFOCOM, Proceedings of*, 2015.
- [54] S. Wang, X. Meng, and T. Chen, “Wide-area control of power systems through delayed network communication,” *Control Systems Technology, IEEE Transactions on*, vol. 20, pp. 495–503, Mar. 2012.
- [55] H. Gao, X. Meng, and T. Chen, “Stabilization of networked control systems with a new delay characterization,” *Automatic Control, IEEE Transactions on*, vol. 53, pp. 2142–2148, Oct. 2008.



- [56] A. Agarwal and J. C. Duchi, “Distributed delayed stochastic optimization,” in *Advances in Neural Information Processing Systems*, pp. 873–881, 2011.
- [57] N. Gatsis and G. B. Giannakis, “Residential load control: Distributed scheduling and convergence with lost AMI messages,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 2, pp. 770–786, 2012.
- [58] S. F. Bush, S. Goel, and G. Simard, “IEEE vision for smart grid communications: 2030 and beyond,” *IEEE Vision for Smart Grid Communications: 2030 and Beyond*, pp. 1–390, May 2013.
- [59] M. Govindarasu and C.-C. Liu, “Cyber physical security of bulk power system: From fault resilient grid to attack resilient grid,” in *National CPS Energy Workshop*, 2013.
- [60] Electric Power Research Institute, “Electric sector failure scenarios and impact analyses,” tech. rep., Electric Power Research Institute, Sep. 2013.
- [61] American Electric Power System, “118 bus power flow test case,” 1962. [https://www.ee.washington.edu/research/pstca/pf118/pg\\_tc118bus.htm](https://www.ee.washington.edu/research/pstca/pf118/pg_tc118bus.htm).
- [62] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 1999.
- [63] D. Energy, “Duke energy to begin process to exit its midwest generation business,” 2014.
- [64] S. Newell and A. Faruqui, “Dynamic pricing: Potential wholesale market benefits in new york state,” *The Brattle Group*, 2009.
- [65] P. Siano, “Demand response and smart grids—a survey,” *Renewable and Sustainable Energy Reviews*, vol. 30, pp. 461–478, 2014.
- [66] J. Hansen, J. Knudsen, A. Kiani, A. M. Annaswamy, and J. Stoustrup, “A dynamic market mechanism for markets with shiftable demand response,” in *IFAC World Congress, Proceedings of*, 2014.
- [67] P. K. Bitar, Eilyan and P. Varaiya, “Coordinated aggregation of distributed demand-side resources,” *Power Systems Engineering Research Center Publication 14-12*, 2014.
- [68] S. Granville, J. Mello, and A. Melo, “Application of interior point methods to power flow unsolvability,” *Power Systems, IEEE Transactions on*, vol. 11, no. 2, pp. 1096–1103, 1996.
- [69] D. Hammerstrom, R. Ambrosio, J. Brous, T. Carlon, D. Chassin, J. DeSteese, R. Guttromson, G. Horst, O. Järvegren, R. Kajfasz, *et al.*, “Pacific northwest gridwise testbed demonstration projects,” *Part I. Olympic Peninsula Project*, 2007.
- [70] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence properties of the nelder–mead simplex method in low dimensions,” *SIAM Journal on optimization*, vol. 9, no. 1, pp. 112–147, 1998.
- [71] R. R. Barton and J. S. Ivey Jr, “Nelder-mead simplex modifications for simulation optimization,” *Management Science*, vol. 42, no. 7, pp. 954–973, 1996.

- [72] S. Katipamula, D. P. Chassin, D. D. Hatley, R. G. Pratt, and D. J. Hammerstrom, "Transactive controls: Market-based grid-wisely controls for building systems," *Pacific Northwest National Laboratory, Richland, WA.* [Online]. Available: [http://www.pnl.gov/main/publications/external/technical\\_reports/PNNL-15921.pdf](http://www.pnl.gov/main/publications/external/technical_reports/PNNL-15921.pdf), 2006.
- [73] L. Chen, N. Li, S. Low, and J. Doyle, "Two market models for demand response in power networks," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pp. 397–402, Oct 2010.
- [74] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 320–331, 2010.
- [75] A. Agarwal, O. Chapelle, M. Dudík, and J. Langford, "A reliable effective terascale linear learning system," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1111–1133, 2014.
- [76] A. Nedić and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *Automatic Control, IEEE Transactions on*, vol. 54, no. 1, pp. 48–61, 2009.
- [77] P. Wood, S. Bagchi, and A. Hussain, "Defending against strategic adversaries in dynamic pricing markets for smart grids," in *2016 8th International Conference on Communication Systems and Networks (COMSNETS)*, pp. 1–8, Jan 2016.
- [78] R. Tan, V. Badrinath Krishna, D. K. Yau, and Z. Kalbarczyk, "Impact of integrity attacks on real-time pricing in smart grids," in *ACM SIGSAC conference on Computer & communications security, Proceedings of*, pp. 439–450, ACM, 2013.
- [79] C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava, "CPS: market analysis of attacks against demand response in the smart grid," in *Computer Security Applications Conference, Proceedings of*, pp. 136–145, ACM, 2014.
- [80] P. Wood, "Source code for model and simulation," 2015.
- [81] Z. Wang, A. Scaglione, and R. J. Thomas, "Generating statistically correct random topologies for testing smart grid communication and control networks," *Smart Grid, IEEE Transactions on*, vol. 1, no. 1, pp. 28–39, 2010.
- [82] A. Laszka, M. Felegyhazi, and L. Buttyan, "A survey of interdependent information security games," *ACM Computing Surveys (CSUR)*, vol. 47, no. 2, p. 23, 2014.
- [83] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pp. 895–902, International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [84] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordóñez, "Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service," *Interfaces*, vol. 40, no. 4, pp. 267–290, 2010.

- [85] M. Brown, W. B. Haskell, and M. Tambe, "Addressing scalability and robustness in security games with multiple boundedly rational adversaries," in *Decision and Game Theory for Security*, pp. 23–42, Springer, 2014.
- [86] C. Barreto, J. Giraldo, A. Cardenas, E. Mojica-Nava, and N. Quijano, "Control systems for the power grid and their resiliency to attacks," *Security Privacy, IEEE*, vol. 12, pp. 15–23, Nov. 2014.
- [87] NESCOR, "Electric sector failure scenarios and impact analyses," tech. rep., Electric Power Research Institute, Incorporated, 2013.
- [88] C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava, "Cps: Market analysis of attacks against demand response in the smart grid," in *ACSAC '14*, 2014.
- [89] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 1, p. 13, 2011.
- [90] Nyiso, "6/22/2015-6/23/2015 nyc lbmp," 2015. [http://www.nyiso.com/public/markets\\_operations/market\\_data/custom\\_report/index.jsp?report=rt\\_lbmp\\_zonal](http://www.nyiso.com/public/markets_operations/market_data/custom_report/index.jsp?report=rt_lbmp_zonal).
- [91] J. N. Russell Hensley and M. Rogers, "Battery technology charges ahead," *McKinsey and Company*, 2012. [http://www.mckinsey.com/insights/energy\\_resources\\_materials/battery\\_technology\\_charges\\_ahead](http://www.mckinsey.com/insights/energy_resources_materials/battery_technology_charges_ahead).
- [92] M. Krotofil, A. A. Cárdenas, B. Manning, and J. Larsen, "CPS: driving cyber-physical systems to unsafe operating conditions by timing DoS attacks on sensor signals," in *Proceedings of the 30th Annual Computer Security Applications Conference*, 2014.
- [93] J. J. Santanna, R. Durban, A. Sperotto, and A. Pras, "Inside booters: An analysis on operational databases," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pp. 432–440, May 2015.
- [94] J. Wang, M. Shahidehpour, and Z. Li, "Security-constrained unit commitment with volatile wind power generation," *Power Systems, IEEE Transactions on*, vol. 23, no. 3, pp. 1319–1327, 2008.
- [95] M. Roozbehani, M. Dahleh, and S. Mitter, "Dynamic pricing and stabilization of supply and demand in modern electric power grids," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pp. 543–548, IEEE, 2010.
- [96] P. Wood, C. Gutierrez, and S. Bagchi, "Denial of service elusion (dose): Keeping clients connected for less," in *Reliable Distributed Systems (SRDS), 2015 IEEE 34th Symposium on*, pp. 94–103, Sept 2015.
- [97] A. Nixon and C. Camejo, "Ddos protection bypass techniques," *Black Hat USA*, 2013.
- [98] M. E. Donner, "Increasing size of individual ddos attack define third quarter," Sept. 2013.

- [99] R. Masse, “Denial of service as a service - asymmetrical warfare at its finest,” Dec. 2013.
- [100] S. Communication, “Ddos protection - secureport global,” Sept. 2013.
- [101] LiquidWeb, “Ddos prevention pricing structure,” Apr. 2014.
- [102] I. CloudFlare, “Cloudflare - compare plans,” Oct. 2013.
- [103] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu, “Portcullis: protecting connection setup from denial-of-capability attacks,” in *SIGCOMM '07*, SIGCOMM '07, (New York, NY, USA), pp. 289–300, ACM, 2007.
- [104] V. Kambhampati, C. Papadopolous, and D. Massey, “Epiphany: A location hiding architecture for protecting critical services from ddos attacks,” in *Dependable Systems and Networks (DSN), 2012 42nd Annual IEEE/IFIP International Conference on*, pp. 1–12, 2012.
- [105] Q. Jia, K. Sun, and A. Stavrou, “Motag: Moving target defense against internet denial of service attacks,” in *ICCCN, 2013*, pp. 1–9, 2013.
- [106] Q. Jia, H. Wang, D. Fleck, F. Li, A. Stavrou, and W. Powell, “Catch me if you can: A cloud-enabled ddos defense,” in *Dependable Systems and Networks (DSN), 2014 44th Annual IEEE/IFIP International Conference on*, pp. 264–275, IEEE, 2014.
- [107] A. D. Keromytis, V. Misra, and D. Rubenstein, “Sos: Secure overlay services,” *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 61–72, 2002.
- [108] A. Stavrou and A. D. Keromytis, “Countering dos attacks with stateless multipath overlays,” in *Proceedings of the 12th ACM conference on Computer and communications security*, pp. 249–259, ACM, 2005.
- [109] D. G. Andersen *et al.*, “Mayday: Distributed filtering for internet services,” in *USENIX Symposium on Internet Technologies and Systems*, pp. 20–30, 2003.
- [110] G. Loukas and G. Öke, “Protection against denial of service attacks: a survey,” *The Computer Journal*, vol. 53, no. 7, pp. 1020–1037, 2010.
- [111] S. Agarwal, T. Dawson, and C. Tryfonas, “Ddos mitigation via regional cleaning centers,” tech. rep., Sprint ATL Research Report RR04-ATL-013177, 2003.
- [112] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, “Feedback control of dynamics systems,” *Addison-Wesley, Reading, MA*, 1994.
- [113] J. Mirkovic, A. Hussain, B. Wilson, S. Fahmy, P. Reiher, R. Thomas, W.-M. Yao, and S. Schwab, “Towards user-centric metrics for denial-of-service measurement,” in *Proceedings of the 2007 workshop on Experimental computer science*, p. 8, ACM, 2007.
- [114] J. Mirkovic, S. Fahmy, P. Reiher, and R. K. Thomas, “How to test dos defenses,” in *Conference For Homeland Security, 2009. CATCH'09. Cybersecurity Applications & Technology*, pp. 103–117, IEEE, 2009.

- [115] J. Mirkovic, A. Hussain, S. Fahmy, P. Reiher, and R. Thomas, “Accurately measuring denial of service in simulation and testbed experiments,” *Dependable and Secure Computing, IEEE Transactions on*, vol. 6, no. 2, pp. 81–95, 2009.
- [116] S. Triukose, Z. Wen, and M. Rabinovich, “Measuring a commercial content delivery network,” in *Proceedings of the 20th international conference on World wide web*, pp. 467–476, ACM, 2011.
- [117] E. Nygren, R. K. Sitaraman, and J. Sun, “The akamai network: a platform for high-performance internet applications,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, 2010.
- [118] P. Wood, “Agent-based model for dose.”
- [119] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker, “Ddos defense by offense,” in *SIGCOMM '06*, SIGCOMM '06, (New York, NY, USA), pp. 303–314, ACM, 2006.
- [120] A. W. Services, “Amazon ec2 pricing, pay as you go for cloud computing services,” Sept. 2013.

VITA

## VITA

Paul Wood received his Bachelor of Science (BS) degree in Electrical Engineering in 2010 from Tennessee Technological University, in Tennessee, USA. He joined Purdue's Dependable Computing Systems Lab (DCSL) in the Fall of 2011, and began work on his Ph.D under Prof. Saurabh Bagchi. His interests lie in control systems, large computing system architecture, dependability, and networking.