Optimizing Defensive Investments in Energy-Based Cyber-Physical Systems

Paul Wood
and Saurabh Bagchi
School of Electrical and Computer Engineering
Purdue University
West Lafayette, Indiana, USA
Email: [pwood,sbagchi]@purdue.edu

Alefiya Hussain
ISI/USC
Marina del Rey, California, USA
Email: hussain@isi.edu

***Abstract*—Interdependent cyber-physical systems (CPS) connect physical resources between independently motivated actors who seek to maximize profits while providing physical services to consumers. Cyberattacks in seemingly distant parts of these systems have local consequences, and techniques are needed to analyze and optimize defensive costs in the face of increasing cyber threats. This paper presents a technique for transforming physical interconnections between independent actors into a dependency analysis that can be applied to find optimal defensive investment strategies to protect assets from financially motivated adversaries in electric power grids.**

***Keywords*-cyber-physical attacks; cyber-physical systems; smart grid; computer security; power system security; game theory**

I. Introduction

Industrial control systems (ICS) which drive cyber-physical systems (CPS) are becoming more interconnected throughout all domains and across corporations. Successful attacks on CPS's are becoming more visible [1], as demonstrated by Stuxnet [2] and shown in a recent ICS-CERT report [3]. Strongly networked and interconnected embedded systems are the cornerstone of most CPS's, providing and coordinating distributed controls often across wide area networks. Protecting all of these systems, however, is impractical–finite resources are allocated to security. The practitioner community 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 interconnected companies face attacks. With such a framework, risk-aware defensive strategies can be formulated for minimizing the impact of security attacks.

Understanding and measuring the complex interactions that occur in interdependent CPS and creating an optimal response to attacks is an important concern in today's world of rising security threats. In a CPS such as the electric power grid, 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 their fates are tied in complex ways. Enron demonstrated in the 2000 California Power Crisis [4] that carefully placed outages could net huge profits. If we map these manipulations to perturbations that a dedicated adversary could introduce in the system, then it becomes clear that we need a scientific basis for seeing which components are most vulnerable and where defensive investments are likely to have the best outcome.

This paper addresses the problem of asset protection in the face of strategic adversaries in electric power grids. The model that we develop in the paper 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 generators(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 to secure the cyber systems utilized by the asset. We explore the gamut of relationships that can exist between the actors as they relate to the defensive strategies that they deploy. 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 financially motivated, then attacks are driven by profits and defenses are driven by losses. This allows 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 that 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 that 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.

Existing techniques in [5–9] have focused on optimizing and planning resource allocation under constraints that primarily prevent shortages while minimizing costs, but these techniques do not consider the interdependent nature of decisions made by entities in an interconnected CPS along with their security implications. Prior development [10–15] creates a foundation for solving some of the larger defensive decisions problems, but it does not consider the interdependence between the CPS and the multi-objective nature of an intelligent, profit-motivated attacker and distributed defenders. An impact analysis technique is needed as the foundation for a utility function in a larger framework that captures the unique characteristics seen in interdependent CPS.

Section II provides an overview of the solution with Section II-D describing the impact model, Section II-E describing the strategic adversary, and Section II-F detailing how the defenders react. Section III contains the experimentation, Section IV the related work, and Section V concludes the paper.

II. The Model

This section provides an overview of our economic model in energy-based cyber-physical systems.

A. System Goals

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 that measures the financial outcomes of perturbations in the physical system, which are driven by cyber-attacks. Embedded systems operating in the CPS are attacked resulting in reduced asset productivity. 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.

B. Multiple Actors

A key divergence from prior work in the techniques presented in this paper is the presence of multiple independent actors. In an interdependent CPS, several different companies are competing for revenues and profits in an open market, and these independent actors represent potential benefactors in cases of malice or disruption. When constructing dependable systems, metrics are often driven from a monolithic ownership perspective (total throughput, etc.). In wide-area energy networks, however, the economic impact of each player must be considered.

C. Attacks and Impacts

From the perspective of multiple actors, an adversary can create disruptions in the system in a way that profits some actors while hurting the system overall. Attacks in this scenario extend as compromising control systems to disrupt physical flows. Specific attack mechanisms such as buffer overflow exploits, etc., are not considered in this paper.

Whenever an attack is launched, the impact can be measured in two ways. First there is a reduction in the efficiency or capacity of the overall system–a more traditional metric that measures pre and post-attack profits, with post-attack profits always being equal or lower to pre-attack levels. Second there is the independent impact on the multiple actors in the system. In this case, the combined change in profit is always negative, but there may be individuals who benefit from the attack. This is the intuition that drives the techniques presented in this paper.

D. The Impact Model

The impact model comes from analysis of energy markets. In these markets, the flow of energy is scheduled with an objective of minimizing costs at producers and maximizing utility at consumers. Techniques such as Optimal Power Flow (OPF) [16] assign power flows to maximize Social Welfare (SW). In the adaptation used in this paper, a graph is created to capture the power flow structure for a medium-term window of time, ignoring the low level mechanics such as voltages and phase angles required to achieve the required flows. The following section describes how a graph structure is created and optimized to reach a maximal Social Welfare.

In an energy-based CPS, the producers have a per-unit cost and the consumers a per-unit price for energy which

drives strategies for energy flow. The CPS (and supporting market mechanisms) optimizes the costs and revenues to provide for the cheapest energy flows to the highest paying customers. Traditionally these markets are slow, updating every 15 minutes, but because electric energy cannot be easily stored interconnected embedded systems are becoming active market players [17]. Consequently, the energy flows in energy CPS are becoming much more dynamic–outages and changes in demand can result in large and localized price levels.

1) Social Welfare: The basic energy-based CPS is modelled as a flow graph. The hubs or vertices serve as electrical buses or gas pipe headers and allow energy to flow via edges from sources to sinks. Each edge has a capacity, loss, and cost associated with the flow of energy. For generalizable equations, the cost may be negative to represent revenues. The fundamental optimization problem is determining the flow levels through each edge in the graph. Sources with cheap energy costs are likely to have high flows. Similarly, consumers with high revenues will see the most energy. For the sake of simplicity in algorithmic convergence, the per-unit costs are assumed fixed at the producers and consumers.

The concept of social welfare or utility follows system-wide profitability. Maximum revenues and minimum costs provide for the most social welfare–if a single company owned all assets this is the best decision to make. Linear programming is used to solve for the optimal flows in the system, and utility function is defined as follows. Each edge has a capacity $c(u, v)$, loss $l(u, v)$, and cost $a(u, v)$. Table I contains the variable and function names for reference. The profits are maximized with the listed constraints. Compared to traditional power system optimizations, these constraints do not consider the stability of the grid (generator response time) other than power flow limitations. New technologies (specifically D-FACTS [18]) allow for a more simplified view of grid planning.

$$\text{Utility} = \min \sum_{(u,v)\in E} a(u, v) \cdot f(u, v) \tag{1}$$

Subject to constraints:

$$0 \leq f(u, v) \leq c(u, v) \tag{2}$$

$$d(v) \leq \sum_{u\in V} c(u, v) \text{for all } v \in \text{L} \tag{3}$$

$$s(v) \geq \sum_{u\in V} c(v, u) \text{for all } v \in \text{G} \tag{4}$$

$$\sum_{u\in V} f(u, v) \leq d(v) \text{ for all } v \in \text{L} \tag{5}$$

$$\sum_{v\in V} f(u, v) \leq s(u) \text{ for all } u \in \text{G} \tag{6}$$

Table I
LIST OF PARAMETER AND FUNCTION DESCRIPTIONS

a(u,v)	Unit cost from u to v
c(u,v)	Capacity
d(v)	Demand
s(v)	Supply
f(u,v)	Actual flow
l(u,v)	Loss percentage
L	Set of all sinks/loads
G	Set of all sources/generators
U	Utility
I	Impact
P_a	Probability of Attack
P_s	Probability of Success, Given Attacked
$C_{dt}(t)$	Cost of Defending Target t
$C_{atk}(t)$	Cost of Attacking Target t

$$\sum_{w\in V} \frac{f(u, w)}{1 - l(u, w)} = \sum_{w\in V} f(w, u) \ \forall \ u \tag{7}$$

The first equation 1 measures profit as a function of cost and flow in the system. Equations 3 and 4 constrain demand to levels desired or possible in the system, and equation 2 enforces capacity constraints. Equations 5 and 6 serve to prevent over-production or over-selling. Equation 7 conserves energy flows through the vertices in the graph. The division by $1-l(u, w)$ accounts for losses in transmission by requiring more total input than output from hubs whenever energy is being transmitted.

2) Social Welfare with Multiple Actors: The utility function listed above provides optimal social welfare. A problem arises, however, when the society is composed of multiple independent actors who are competing for profits. In this case, the optimal flows cannot be decided from a single point of interest. This section describes an algorithm for determining individual utilities when multiple actors exist in the system.

Determining how independent players will negotiate on prices is complex and often modeled in game theory. Work in [19] for example analyzes games that generate prices for buyers and sellers in completely open energy markets using interconnected smart meters. The focus in this work, however, is not in analyzing these particular games but rather in providing a reasonable estimate under some assumptions. The first assumption is that flows always operate at the social welfare optimum. Intuitively, if the actors were willing to cooperate with each other, then this equilibrium would always be reached (as a collation-proof Nash equilibrium). The individual utilities then are not found by optimizing flows. Instead, only the profits of the system must be distributed.

The profits are divided among the actors in the system by evaluating marginal costs at each point in the system.

Notionally, the marginal cost represents the price of the alternative (i.e. competition). Since the actors are independent, they will not collaborate with each other to influence prices. Therefore, all competition is assumed perfect in the sense that each actor will charge the maximum up to the marginal cost at which point a competitor will overtake production. The marginal cost is calculated by fixing the flows for each actor by adding a constraint to the optimization problem and reducing the capacity of each positive-flow edge by one unit. The reduction in utility is the corresponding marginal cost. This is done independently for each actor and each edge.

```
1: For each actor, fix edge outflows
       and determine marginal cost
2: Assign fraction of marginal cost to
    edges in single actor problem
3: Increase fraction until flow is
       perturbed
4: Reduce cost at flow perturbed
       edges until flow is restored
       fixing cost at those edges
```

While the marginal cost provides a basis for competition, there are a few profit sharing scenarios that must be resolved algorithmically. In the case of competitors in series, for example, each would arrive at the same marginal cost on its output. They are effectively forced to collaborate on flow levels as a function of physical layout. The middle actor cannot resell what it does not receive. To rectify this, the profit taken by each actor in series is incrementally grown as listed below. Consequently, the actors receive a portion of the profit roughly equal to $1/N$.

```
1: Set the supply/demand to its maximum
   value at the interface between actors
2: Optimize local profits
3: Pass the actual supply used by each
   actor to adjust the demand at the
   interfacing actors
4: Repeat 1 3 for each actor until d(u)
   converges within a tolerance (0.5 %)
```

3) Impact Measurement and Attacks: The impact is measured by perturbing some part of the model (cost, loss, etc.) as follows. Impact $=$ Utility$'$ $-$ Utility where Utility$'$ $=$ Utility as $a, c, l \rightarrow a', c', l'$.

Attacks in the model are directly represented by augmenting the different model parameters (effectively changing the graph itself). In a realistic scenario, the adversary would compromise a control system via an advance persistent threat or other control system vulnerability and obtain an ability to influence the system. The attack could be abrupt and eliminate the capacity of a generator or transmission line. It could also be more subtle and cause slight increases in loss, for example.

4) Knowledge Perturbations: Similar to an actual attack, the model may also be perturbed to represent uncertainty about parameters in the system. An adversary, for example, may collect system information from public sources or via inspection (e.g. from a satellite photograph). To model this behavior, each parameter in the system is perturbed by a normal distribution with a mean centered at the original value. A parameter σ represents the knowledge level as a normal distribution i.e. $c'(u, v) = \mathcal{N}(c(u, v), \sigma^2)$.

5) Model Limitations: Often in grid planning, day-ahead projections are used to schedule around generator constraints. For example, it may take several minutes (or hours) for generating facilities to achieve maximum output. While this is an important concept in long term planning, the flows presented here capture only particular time intervals where demand and generation are expected to be fixed and attainable. A time-domain component can be added to the model by integrating several instances of the utility function to represent varying demands and generating constraints. The approaches presented in this paper, however, are designed and evaluated only for a single demand instance that is assumed to extend for the duration of an attack.

E. Profiting from Manipulations

Perturbations in the model result in changed profits for each actor, and a strategic adversary (SA) attempts to cause perturbations in a profitable way. Disrupted flows will cause losses for the owner of an attacked asset while changing market conditions for other players in the system. For example, an outage may result in a competitor elimination type of a scenario allowing additional profits to be extracted by certain players in the system. In other cases, such as a consumer distribution outage, every actor in the system loses from reduced sources of revenue. Based on this intuition, a strategic adversary can profit from the system by finding targets that benefit a subset of actors with whom she has a vested interest.

1) Profit Collection: In the case of a perturbation (successful attack), the SA is able to collect some percentage of profit or loss that an actor experiences. This is generally achieved by buying and selling stocks or energy delivery futures. For example, the SA might buy stock in actor A, launch an attack which profits A causing the stock to rise in value, and then selling the stock to recover a profit. Similarly, the SA may buy a future delivery contract for energy at a pre-negotiated price, perturb the system, and the resell the energy at an elevated (or reduced) price for a profit or loss. In this way, the SA is able to collect profits in the system.

2) Attack Vectors: Each edge in the graph represents a physical component or asset in the energy system. Each asset (or target from the SA's perspective) may be controlled by a host of control system devices. These devices are not homogeneous, and some effort must be expended for attacking each target. For example, a SA may launch an advance persistent threat that requires careful design, reconnaissance, and re-design of the viruses to successfully disrupt the physical system. Therefore, the cost to attack each target is not fixed and must be evaluated by the adversary to

determine the best targets to attack. More generally, the attack vector chosen has a cost and potential return, and the SA must optimize this decision process.

3) Selecting Targets: The SA performs the following optimization problem using mixed integer linear programming (MILP). The value $IM[a,t]$ is the impact matrix that measures the profit from perturbing target t on actor a. If the value is negative, then that actor experiences a loss for that target. Each target $t \in T$ has an expected cost of attack $C_{atk}(t)$ 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 optimizes her returns as follows:

$$\max_{T,A} \sum_{i \in T} \left(-C_{atk}(i) + \sum_{j \in A} IM[j,i] \cdot T(i) \cdot A(j) \cdot P_s(i) \right) \tag{8}$$

Subject to constraints:

$$T(i) \in \{0,1\} \tag{9}$$

$$A(j) \in \{0,1\} \tag{10}$$

$$\sum_{i \in T} (T(i) \cdot C_{atk}(i)) \leq M_A \tag{11}$$

Equation 8 selects A and T to maximize the return on investment (ROI). Equation 9 and 10 force integer values for target/actor selection, and equation 11 enforces budget constraints. The value of a target is approximated as linearly additive since a single attacker is considered, though some choices may be submodular or supermodular. The set A is not fixed or constrained, and if A is every actor, the target set T will be empty because the underlying system is operating at a maximal social welfare.

4) Limitations of the Adversary Model: While the SA model presented captures the adversary's operational characteristics, estimating the probability of successful attack and the cost to execute an attack can be difficult from the adversary's perspective. Some gentle probing, however, can provide first level approximation. The point of these parameters is to allow for exploration of defense methodologies in the sense that adding layers of security reduces the probability of successful attack and increase the cost of an attack. The SA model can become computationally difficult to solve as the system grows in both the number of actors and targets. This problem can be alleviated to some extent by partitioning the system and actors into a divide-and-conquer algorithm. The submodular and supermodular concerns can be alleviated by limiting the attacker's budget as a complex multi-system attack is rather unlikely.

F. 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 a 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.

1) 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:

$$\max_{D} \sum_{t \in T_a} (P_a(t) \cdot I(a,t) \cdot (1 - D(t)) - C_d(t) \cdot D(t)) \tag{12}$$

Subject to the constraint:

$$D(t) \in \{0,1\} \tag{13}$$

$$\sum_{t \in T_a} (D(t) \cdot C_d(t)) \leq M_D(a) \tag{14}$$

Equation 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 14 which caps the amount of expenditures on defense to M_D. This can be solved using MILP, as in the strategic adversary case.

2) 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''.

3) 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:

$$C_c d(a,t) = \frac{C_d(t) \cdot I(a,t)}{\sum_{i \in CD(t)} I(i,t)} \tag{15}$$

Optimize:

$$\max_D \sum_{i \in T} \left(\sum_{j \in CD(i)} (P_a(j,i) \cdot IM[j,i] \cdot (1 - D(i))) - C_d(i) \cdot D(i) \right) \tag{16}$$

Subject to the constraints:

$$D(i) \in \{0,1\} \tag{17}$$

$$\sum_{i \in T_a} (D(i) \cdot C_c d(j,i)) \leq M_D(j) \ \forall j \in A \tag{18}$$

These equations are identical to the earlier set when $|CD(t)| = 1 \ \forall t$. The optimization in Equation 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.

4) Defender Model Discussion: The defender model provides a basis for evaluating dependability from the perspective of an impending strategic adversary. In a traditional dependability model, the defender evaluates self-loss and proportionally protects assets as a way to mitigate those losses. In the SA model the defender is performing the same analysis, but the likelihood of those losses are driven by a for-profit adversary. In this way, traditional dependability models can be augmented with probability of failures that include security-oriented attack probabilities.

III. Experimentation

A. CPS Model

1) Physical Model: An interdependent gas-electric system comprised of six western US states was captured for experimentation. A simulation was constructed using the algorithms described in the paper in MATLAB/Octave using "linprog" and the GLPK linear programming solvers. The model is based on information available from the Energy Information Administration (EIA) [20, 21]. Each state's vertex in the graph corresponds with its geographic centroid, for purposes of calculating per-unit transmission losses. Each state has two vertices with two consumers–one for gas and one for electricity. In total there are 12 vertices and 18 long haul transmission edges. Figure 1 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).

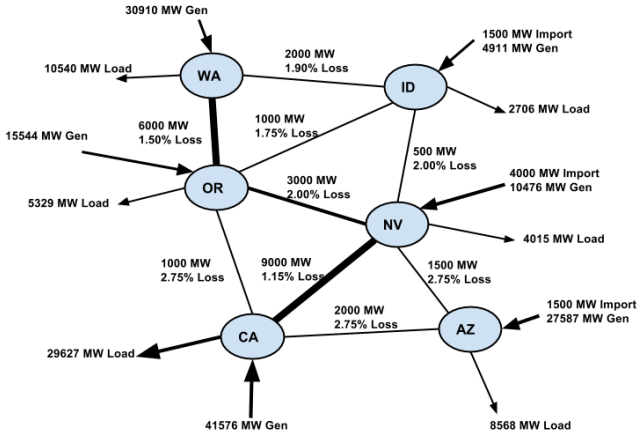

(a) Electric Model

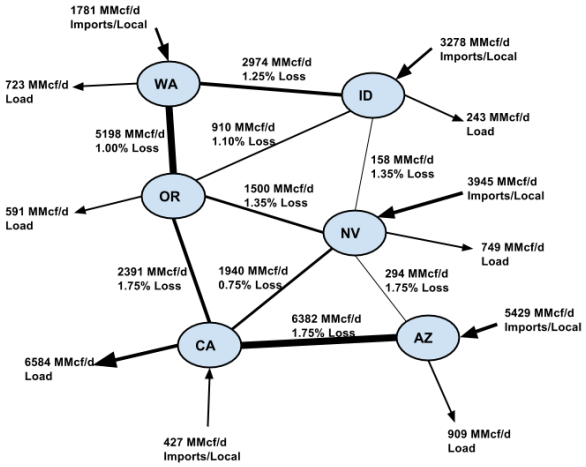

(b) Natural Gas Model

Figure 1. A flow model is created for six Western US states for both an electric (a) and natural gas (b) infrastructure.

2) 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, allowing for transportation costs. For the loss function l, a calculation is made based on a typical loss of 1% per 400 km [22], since the actual loss rates vary based on each individual pipeline's construction characteristics. The resulting loss rates are seen in Figure 1 (b). The capacity function c directly maps to EIA's dataset [20]. For the loss on the gas to electricity transformation, we use each state's energy information profile. Finally the supply, imports, and demand for each state were calculated by converting yearly consumption into smaller time scale amounts. The values for the electric infrastructure are calculated similarly.

Similarly, the values for the electric infrastructure are calculated using the EIA sources [21]. 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 [23].

We wish to create a more challenging model to evaluate noticeable impact of an attack (and the resultant defensive investments). To create a more challenging model, we make several modifications to the baseline model presented above. 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.

3) Attacker and Ownership Model: The profit-seeking SA may select any attack vector and model perturbation in practice. In this experiment, the perturbation for a particular target is to reduce its capacity to zero, modeling an outage scenario. This could easily be achieved by crashing a programmable logic controller (PLC) rather indiscriminately, and is thus a moderate complexity attack.

For each experiment, multiple random sets of actors are created and measured, and the results taken as means across these variations. The distribution is that if there are N actors, each asset has a $frac1N$ chance of belonging to any particular actor.

B. 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 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.

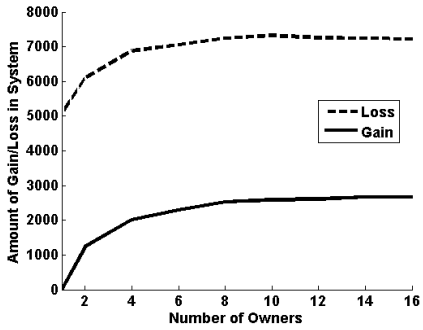

Figure 2. The total gain and loss in the system, as the sum across impacts felt by all actors, increase as the number of actors in the system increase up to a point of saturation. The sum of the gain and negative loss remain constant.

C. 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 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 II-E3. 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 (σ) 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 3 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. As the knowledge level of the attacker is decreased, the effectiveness of the attack also decreases due to poorer decision making.

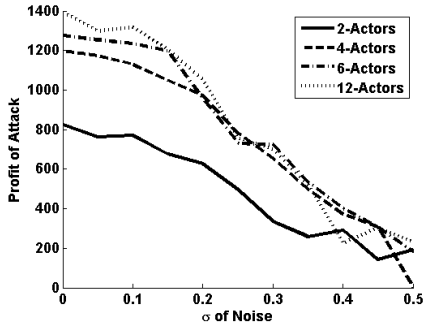

Figure 3. This figure shows the profitability of the strategic adversary versus the amount of knowledge (inverse of noise) that it has about the system. As the noise increases, the profitability decreases. Additionally, as the number of actors increases, the profitability of the SA also increases because of profit opportunities.

Figure 4 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 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).

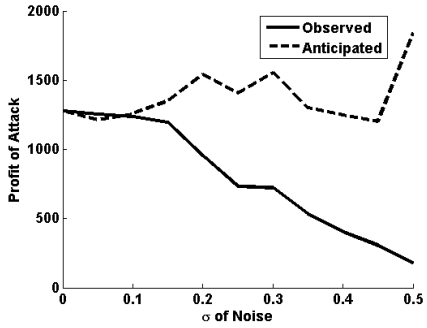

Figure 4. This compares the profit of attack for a 6-actor system. The SA anticipated returns, based on the noisy model, do not decay with knowledge level. This means that if the SA is overconfident, the observed returns will be much less than anticipated.

D. 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.

The defender's goal is to minimize the impact of an attack. The metric we use for this experiment is then 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 that cause a significant loss to itself and are likely to be attacked. This it does under incomplete knowledge (hence the σ 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 σ for the various parameters that the adversary uses). This mechanism is as detailed in Section II-F2.

Figure 5 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.

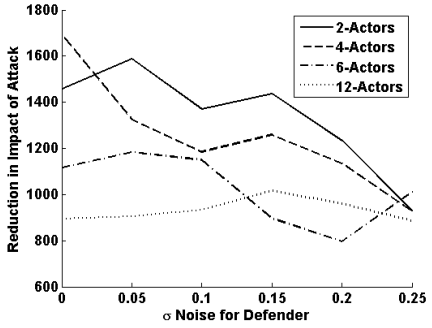

Figure 5. The effectiveness of a defense is graded by its impact reduction in ground truth versus the knowledge level of the defender, modeled as noise added to the ground truth. As the number of actors increases, the effectiveness of the defense decreases due to misaligned incentives and a lack of pooled defensive budgets.

Figure 6 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 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.

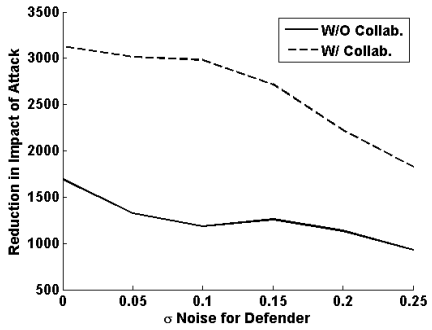

Figure 6. The impact of collaboration is measured by allowing the actors to share in defensive costs. When the costs are shared, more effective investments can be made.

Figure 7 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 5 that the effectiveness of defense decreases.

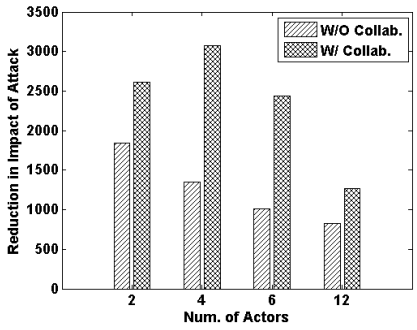

Figure 7. Collaboration allows actors to improve their defenses. In this case, the system-wide defensive investment is fixed as the number of actors increases, resulting in reduced benefit of collaboration as the number of actors increases and their individual budgets dwindle.

IV. Related Work

A collection of games surveyed in [24, 25], called interdependent information security games, evaluate the impact of defensive decisions made by one player on the others. These games define the players' interaction generally in a contagion-type model[26, 27], whereas this paper focuses on interactions that occur as a result of physical interconnection. Some games target network control systems specifically [28], however these models do no address the physical interdependent [25, 29] aspect of security.

The finances related to attacks have been studied in [29]. The game provides for reduced returns on investments by the attacker should the defender make a defensive move for the right asset. This model is useful for evaluating how a strategic adversary might be impacted in the long term. In this paper, however, the short term impacts are studied.

A. Electric System Modeling

Solving for optimal energy flows in interconnected energy systems is not a new problem, and many solutions exist for providing and optimizing flows around well-developed constraints. Most of these techniques, however, are not suitable for multi-player games because a single player (the independent system operator or ISO) is responsible for collecting and managing bids in a closed market. This work focuses on the market aspects more than the constraint aspects as in security constrained unit commitment (SCUC) [5] which provides system stability (security) during failures. The SCUC method has been extended to interconnected gas-electric systems and provides for useful, however complex planning [6–9]. This paper focuses on the multi-player aspects once a particular energy flow has been established by the simple flow-optimization problem provided.

B. Computer Security and Graphs

General approaches to computer security are not able to map to the objectives (profit) that this paper focuses on. Information stolen in computer hacks is discrete, non-time varying, and thus generally binary. Work in [30, 31] identifies at-risk components using graphical interdependencies with this binary attack objective in mind. Using pure graphic methods in determining the risk of components was done in [32], but viability of the approach was questioned in other research [33]. Existing graph-based techniques are useful for establishing some level of interdependency [34], but the problem of developing a security strategy around physical flow perturbations has not been adequately addressed for interdependent CPS.

C. Games and Defense

Several games have been studied when the impact of attack known outside of a CPS context. Several recent techniques [11–15] evaluate how game theory can apply to different grid-based scenarios. This work supplements these

contributions by providing a strategic adversary model with multiple profit-impacted actors. More research is needed to evaluate the impacts of interdependencies in the physical domain with defender and attacker strategies.

V. Conclusion

In this paper, we present a modeling technique for evaluating cybersecurity defensive investments in interconnected cyber-physical systems. An impact analysis technique enables multiple actors to compete and maximize their individual profits in a flow-optimization problem. The multi-actor approach allows 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. We find 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.

References

[1] 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, no. 4, pp. 853–865, Jul. 2010. [Online]. Available: http://dx.doi.org/10.1109/TSMCA.2010.2048028

[2] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *Security & Privacy, IEEE*, vol. 9, no. 3, pp. 49–51, 2011.

[3] ICS-CERT, "Ics-cert monitor, internet accessible control systems at risk," 2014.

[4] T. Egan, "Tapes show enron arranged plant shutdown," 2005.

[5] 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*. IEEE, 2006, pp. 1759–1766.

[6] 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.

[7] 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*. IEEE, 2007, pp. 468–472.

[8] 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.

[9] 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.

[10] R. Gopalakrishnan, J. R. Marden, and A. Wierman, "An architectural view of game theoretic control," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 3, pp. 31–36, 2011.

[11] 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.

[12] K. J. Ross, "Application of game theory to improve the defense of the smart grid," DTIC Document, Tech. Rep., 2012.

[13] 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.

[14] 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.

[15] 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.

[16] H. W. Dommel and W. F. Tinney, "Optimal power flow solutions," *power apparatus and systems, IEEE transactions on*, no. 10, pp. 1866–1876, 1968.

[17] M. H. Albadi and E. El-Saadany, "A summary of demand response in electricity markets," *Electric Power Systems Research*, vol. 78, no. 11, pp. 1989–1996, 2008.

[18] K. Rogers and T. Overbye, "Power flow control with distributed flexible ac transmission system (d-facts) devices," in *North American Power Symposium (NAPS), 2009*, Oct 2009, pp. 1–6.

[19] A. Kiani and A. Annaswamy, "Wholesale energy market in a smart grid: Dynamic modeling and stability," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 2202–2207.

[20] U. E. I. Administration, "Natural gas," 2014. [Online]. Available: http://www.eia.gov/naturalgas/

[21] ——, "Electricity," 2014. [Online]. Available: http://www.eia.gov/electricity/

[22] U. F. E. R. Commission, "Ferc: Natural gas," 2014. [Online]. Available: http://www.ferc.gov/industries/gas/gen-info/fastr/index.asp

[23] J. M. Petrash, "Long-term natural gas contracts: Dead, dying, or merely resting," *Energy LJ*, vol. 27, p. 545, 2006.

[24] A. Laszka, M. Felegyhazi, and L. Buttyan, "A survey of interdependent information security games," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 23:1–23:38, Aug. 2014.

[25] M. H. Manshaei, Q. Zhu, T. Alpcan, T. Bacşar, and J.-P. Hubaux, "Game theory meets network security and privacy," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, p. 25, 2013.

[26] H. Chan, M. Ceyko, and L. E. Ortiz, "Interdependent defense games: Modeling interdependent security under deliberate attacks," *arXiv preprint arXiv:1210.4838*, 2012.

[27] V. M. Bier, "Choosing what to protect," *Risk Analysis*, vol. 27, no. 3, pp. 607–620, 2007.

[28] 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.

[29] 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.

[30] L. Wang, S. Noel, and S. Jajodia, "Minimum-cost network hardening using attack graphs," *Computer Communications*, vol. 29, no. 18, pp. 3812–3824, 2006.

[31] 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*. ACM, 2007, pp. 204–213.

[32] 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.

[33] 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.

[34] 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.