Integrity Attacks on Real-Time Pricing in Electric Power Grids

RUI TAN, Advanced Digital Sciences Center, Illinois at Singapore
VARUN BADRINATH KRISHNA, Advanced Digital Sciences Center, Illinois at Singapore; University of Illinois at Urbana-Champaign
DAVID K. Y. YAU, Advanced Digital Sciences Center, Illinois at Singapore; Singapore University of Technology and Design
ZBIGNIEW KALBARCZYK, University of Illinois at Urbana-Champaign

Modern information and communication technologies used by electric power grids are subject to cybersecurity threats. This paper studies the impact of integrity attacks on real-time pricing (RTP), an emerging feature of advanced power grids that can improve system efficiency. Recent studies have shown that RTP creates a closed loop formed by the mutually dependent real-time price signals and price-taking demand. Such a closed loop can be exploited by an adversary whose objective is to destabilize the pricing system. Specifically, small malicious modifications to the price signals can be iteratively amplified by the closed loop, causing highly volatile price, fluctuating power demand, and increased system operating cost. This paper adopts a control-theoretic approach to deriving the fundamental conditions of RTP stability under basic demand, supply, and RTP models that characterize the essential behaviors of consumers, suppliers, and system operators, as well as two broad classes of integrity attacks, namely, the *scaling* and *delay* attacks. We show that, under an approximated linear time-invariant formulation, the RTP system is at risk of being destabilized only if the adversary can compromise the price signals advertised to consumers, by either reducing their values in the scaling attack, or by providing old prices to over half of all consumers in the delay attack. The results provide useful guidelines for system operators to analyze the impact of various attack parameters on system stability, so that they may take adequate measures to secure RTP systems.

Categories and Subject Descriptors: B.8.2 [**Performance and Reliability**]: Performance Analysis and Design Aids; K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms: Economics, Security

Additional Key Words and Phrases: Power grid, smart grid, electricity market, real-time pricing, demand response, stability, cyber security

1. INTRODUCTION

Electric power grids are increasingly using modern information and communication technologies (ICTs) to improve system reliability and efficiency. However, these computerized and networking technologies are subject to security threats that range from personal breaches [McLaughlin et al. 2010] to sophisticated cyber attacks launched by hostile organizations to cause widespread outages [The Wall Street Journal 2009]. As a sophisticated cyber-physical system, a power grid features complex closed-loop feed-

A preliminary version of this paper appeared in The 20th ACM Conference on Computer and Communications Security (CCS), 2013.

Author's addresses: R. Tan and D. K. Y. Yau, Advanced Digital Sciences Center, 1 Fusionopolis Way, #08-10 Connexis North Tower, Singapore 138632; V. Badrinath Krishna and Z. Kalbarczyk, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 1308 W. Main Street, Urbana, IL 61801, USA.

This study is supported in part by the research grant for the Human-Centered Cyber-physical Systems Programme at the Advanced Digital Sciences Center from Singapores Agency for Science, Technology and Research (A⋆STAR), and in part by U.S. Department of Energy under award number DE-OE0000097.

DOI: http://dx.doi.org/10.1145/0000000.0000000

back controls in various physical [Grainger and Stevenson 1994] and economic components [Alvarado 1999], which maintain desirable system performance in the presence of dynamics and uncertainties. However, the impacts of cyber attacks against these closed loops on power grids have received limited research attention. Without a systematic understanding of these impacts, system designers and operators will not be able to truly assess how these attacks may undermine the system's ability to provide mission-critical services, and hence take appropriate defensive measures against the possible threats. This paper makes a step in this direction by quantifying, through both analysis and simulations, the impact of cyber attacks on real-time pricing (RTP), an emerging feature of advanced power grids that involves closed-loop controls to stabilize the electricity market.

Dynamic pricing [Barbose et al. 2005] is a widely adopted means to balance electricity generation and consumption. The electricity price in the wholesale market is updated periodically (e.g., every 5 minutes) to match generation with dynamic demand. In contrast, many current retail markets adopt static pricing schemes such as fixed and time-of-use tariffs, under which the consumers have limited incentives to adapt their electricity consumption to market conditions. This lack of incentives results in high peak demands that strain infrastructure capacities and unnecessarily increase operational costs. By relaying the real-time wholesale prices to end customers, RTP has been considered a key feature of the grids of tomorrow, which can reduce over-provisioning and improve system efficiency. In practice, utilities have provided RTP programs to large commercial and industrial customers for more than ten years [Barbose et al. 2005]. Currently, they are increasingly extending RTP programs to small residential customers, as exemplified by Commonwealth Edison Company (ComEd) [ComEd 2014] and Ameren Corporation [Ameren 2014] in Illinois. Moreover, RTP is becoming a legally required option for consumers [Public Act 094-0977 2014].

Unfortunately, as analyzed in [Roozbehani et al. 2012a], there exists a fundamental *information asymmetry* between the system operators and consumers under RTP. Specifically, a system operator needs to determine the price, which is supposed to clear the market, prior to the consumption decisions made by consumers. As the system operator typically has limited knowledge about the consumers, its best practice is to determine the price based on historical demand. As a result, RTP creates a closed loop formed by the mutually dependent real-time prices and price-taking demand [Roozbehani et al. 2012a]. Such a closed loop can increase the system's sensitivity to dynamics and lower its robustness against situational uncertainties. As such, it can be exploited by an adversary whose objective is to destabilize the RTP system. Specifically, small modifications to the signals in the closed loop made by the adversary can be iteratively amplified by the feedback, causing highly volatile price, fluctuating power demand, and increased system operating cost.

In advanced power grids, the real-time price signals are exposed to different security threats at the source, during network transmissions, and at the consumer-level smart meters. Recent studies [McLaughlin et al. 2010; Rouf et al. 2012] have shown that many smart meters lack basic security measures to ensure the integrity and authenticity of their input/output data. Moreover, a Blackhat demonstration [Davis 2009] showed that a worm was able to control 15,000 emulated smart meters in 24 hours. In light of these infrastructure vulnerabilities, imperative questions regarding RTP security include, "Can the malicious compromise of real-time price signals destabilize the system and cause undesirable consequences such as unacceptably high price volatility? If so, to what extent do the price signals need to be compromised?" A main challenge in answering these questions stems from the complex coupling between the attacker actions and the closed-loop RTP system. For instance, an attack to a few smart meters can cause monetary losses to individual victims, but it will not be able to destabilize the whole system. But if the adversary is able to compromise a sufficiently large number of consumers, the pricing mechanisms, which are designed to stabilize the system, may fail to mitigate the attack's impact. This impact may then pervade the whole system due to the iterative feedback. However, it is challenging to quantify these critical stability boundaries accurately, in order to characterize the impact of the attacks.

In this paper, we adopt a control-theoretic approach, which captures the closed-loop nature of the RTP, to deriving fundamental stability conditions under credible integrity attacks. Based on the linearization of general abstract models of supply and demand, the RTP problem is formulated as a classical control problem for a linear time-invariant (LTI) system. We develop a basic pricing algorithm that sets the price adjustment proportional to the observed error between supply and demand. It ensures stability and captures the essence of stability-ensuring RTP systems. Therefore, the security analysis based on this algorithm provides a baseline understanding of the security of these systems. We adopt a control-theoretic metric, namely, the *region of stability*, to characterize the resilience of the closed-loop RTP system with respect to important and practical adversary models. Specifically, we consider two common and broad classes of integrity attacks, which we call the *scaling* and *delay* attacks. In the scaling attacks, the prices advertised to consumers are compromised by a scaling factor, so that the consumers will use scaled prices to make power consumption decisions. In the delay attacks, timing information of prices is corrupted, so that the consumers will use old prices. In addition to directly tampering with data traffic sent to the consumers, these attacks can be accomplished by indirect techniques that are less effort intensive. For instance, the delay attack can be realized by compromising the time synchronization of consumers' smart meters. Note that current commercial smart meters [Schneider Electric 2014] synchronize their clocks by either built-in Global Positioning System (GPS) receivers or a network time protocol (NTP) supported by time servers [Schneider Electric 2014]. Both approaches have been shown to be vulnerable to realistic attack methods [Oracle 2012; Nighswander et al. 2012].

Based on our analytical framework, we derive the region of stability for both the scaling and delay attacks. We show that, under the LTI formulation, the RTP system remains stable if (i) the compromised prices advertised to consumers are amplified versions of the true prices under the scaling attack, or (ii) less than half of the consumers in the pricing system are compromised in the delay attack. On the other hand, if the adversary can break either of these two conditions, the system may experience highly volatile prices and severely fluctuating demand arising from the system instability. To verify our analysis, we conduct several sets of simulations using 1,405 houses, a 4-bus transmission system, and an IEEE 118-bus system, respectively. We demonstrate possible system emergencies (e.g., line overload) and system efficiency degradation (e.g., increased power losses in transmission) caused by the attacks. The insights based on our results are two-fold. First, to destabilize the system and cause system-wide catastrophic consequences, the adversary needs to compromise the price signals to a large number (at least 50% under the delay attack) of the consumers. Although this is indicative of the resilience of RTP systems, the possibility of compromising a large number of smart meters cannot be ignored given known smart meter vulnerabilities and the aforementioned Blackhat demonstration [Davis 2009]. On the other hand, our results suggest that, to achieve the necessary breadth of coverage, the adversary may focus on compromising shared support infrastructures such as NTP time servers. Such attacks are credible as evidenced by existing security incidents such as the Dragonfly viruses [Symantec 2014] that penetrated power grids' ICT systems. Thus, our results highlight the importance of securing these time servers.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the market model. Section 4 defines the RTP stability problem, and develops a control-theoretic formulation of the problem. Sections 5 analyzes the impact of the integrity attacks against the real-time prices for consumers. Section 6 studies the impact of limiting factors in generation on our analytic results. Section 7 discusses other attack models and the impact of various system uncertainties on our analytic results. Section 8 presents simulation results. Section 9 concludes this paper.

2. RELATED WORK

The security of power grids is attracting increasing research attention. In particular, false data injection attacks against the state estimation of power grids have been extensively studied. In [Liu et al. 2011], Liu et al. systematically examine the conditions for bypassing a bad data detection mechanism of state estimation. Later studies [Yuan et al. 2011; Lin et al. 2012; Xie et al. 2011; Jia et al. 2012; Kosut et al. 2011] show that the false data injection attacks can lead to increased system operating costs due to inordinate generation dispatch [Yuan et al. 2011] or energy routing [Lin et al. 2012], as well as economic losses due to misconduct of electricity markets [Xie et al. 2011; Jia et al. 2012; Kosut et al. 2011]. In particular, the studies in [Xie et al. 2011; Jia et al. 2012; Kosut et al. 2011] focus on false data injection attacks on real-time wholesale markets. They primarily focus on the attacks on critical measurements in a power system, which are often well protected by system operators. Moreover, they ignore demand response of end customers to prices. In contrast, we consider integrity attacks that may target distributed smart meters that are much more vulnerable, and also accounts for demand response involving the end customers. Moreover, we focus on how a series of attacks over a time period will affect the pricing system stability. All these related studies [Liu et al. 2011; Yuan et al. 2011; Lin et al. 2012; Xie et al. 2011; Jia et al. 2012; Kosut et al. 2011] analyze attacks on systems using constrained optimization formulations such as economic dispatch. The closed loop characterizing the RTP system in our work imposes specific challenges in the security analysis due to its iterative nature. To address these challenges, we adopt a control-theoretic formulation.

The security of a broader class of cyber-physical systems that feature complex closed loops has been studied recently. In [Cárdenas et al. 2008], Cárdenas et al. identify challenges in the security analysis of these systems. In [Cárdenas et al. 2011], the authors use simulations to study the impacts of integrity and denial-of-service attacks on a chemical reactor with multiple sensors and control loops. In [Amin et al. 2013], the authors perform security threat assessment of supervisory control and data acquisition systems for water supply. These studies focus on demonstrating the possibility of pushing the system to a certain state (e.g., unsafe pressure in a chemical reactor) by tampering with the sensor and/or control signals. In contrast, this paper aims at characterizing the fundamental critical stability conditions of closed-loop RTP systems.

3. MARKET MODEL

This section presents the market model adopted in this paper, which comprises an *independent system operator* (ISO), consumers, and suppliers. This section also presents a basic set of assumptions about the RTP, demand, and supply models. The following sections (Section 4 to Section 8) will introduce necessary variations to these basic assumptions (e.g., linearization and instantiation). Table V in the conclusion section of this paper summarizes these variations. The notation used in this paper and the default physical units of symbols are summarized in Appendix A of the supplementary file containing appendices of this paper.[1] We also use the following mathematical no-

[1] Due to space limitations, all appendices are omitted and can be found in the supplementary file of this paper.

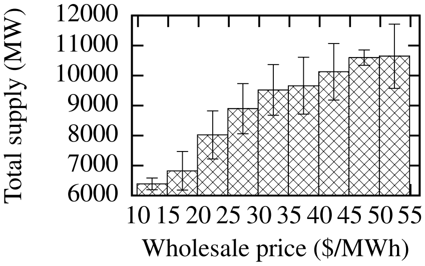

Fig. 1. Total supply vs. wholesale price [Austrilian Energy Market Operator 2014].

tation: $\mathbb{R}^+$/$\mathbb{R}^-$ denotes the set of positive/negative real numbers; $\mathbb{Z}^+$ denotes the set of positive integers; $\dot{f}(x)$ denotes the first derivative of function $f(x)$.

3.1. ISO Model and RTP Schemes

The ISO is a profit-neutral agent, which aims to clear the market, i.e., match supply and demand. It determines a clearing price every T hours and announces it to the suppliers and consumers. Specifically, the price for the kth pricing period $[k \cdot T, (k+1) \cdot T]$, denoted by λ_k, is announced at time instant $k \cdot T$. Hence, this scheme corresponds to ex-ante pricing. We assume that the price must be within a range, i.e., $\lambda_k \in [\lambda_{\min}, \lambda_{\max}]$, where $\lambda_{\max} > \lambda_{\min} \in \mathbb{R}^+$. Note that in many electricity markets, suppliers sell electricity to utilities in wholesale markets, and utilities sell electricity to end consumers in retail markets. The market model adopted in this paper directly relays real-time wholesale prices to end consumers, which preserves the principles of RTP and simplifies the analysis. This model has been employed in previous studies [Roozbehani et al. 2012a] and is consistent with the essence of several experimental RTP programs provided by utilities [Barbose et al. 2005; ComEd 2014; Ameren 2014], which include Board of Public Utilities in New Jersey, Baltimore Gas and Electric Company in Maryland, Duquesne Light in Pennsylvania, ComEd and Ameren in Illinois. In these programs, the hourly wholesale prices published by PJM Interconnection LLC are used directly as retail prices ($T = 1$), where the utilities make profit from fixed service charges only [ComEd 2014; Ameren 2014]. In particular, the ComEd's RTP program, which was developed based on an experimental program [Allcott 2009], had more than 5,000 household participants by 2008 [Allcott 2009] and has been continuously growing in size. A few other RTP programs give customers advance notice of hourly prices. For instance, in the RTP-HA-2 program of Georgia Power [Georgia Power 2014], the price is announced one hour before. To simplify the discussion, we focus on RTP schemes without advance notice. However, our analysis can be easily extended to encompass advance notice. The ISO may apply ex-post price adjustments. In this paper, we assume that the consumption and generation are scheduled according to the ex-ante price only.

In reality, locational prices can be applied to address location-dependent transmission costs. That is, different locations in a transmission system adopt different prices. In this paper, we assume that the transmission lines have high enough capacities such that they are not congested. Under this assumption, as generation cost often dominates transmission cost, variations of prices across different locations at the same moment are small. For instance, according to PJM's operational data [PJM 2014], the relative standard deviation of the locational prices for 219 locations at the same moment is often around 5% only. Note that the price for a location can change significantly over time (e.g., from 10 $/MWh to 55 $/MWh as shown in Fig. 1) . As this paper focuses on the impact of integrity attacks on RTP systems, we ignore the small variations in the locational prices. Thus, we assume that all the suppliers and consumers are subject to the same real-time price at any time instant. Section 8.2 will evaluate the loss of accuracy due to this simplification of assuming identical real-time prices across all the locations. It also discusses the impact of integrity attacks on the power grid when a few transmission lines are already congested before the launch of the attacks.

3.2. Consumers

Let C denote the set of consumers in the system. In the kth pricing period, the demand of a consumer $j \in C$, denoted by $d_j(k, \lambda_k)$, is the sum of the *baseline demand* and *price-responsive demand*, which are denoted by $b_{k,j}$ and $w_j(\lambda_k)$, respectively. By denoting $b_k = \sum_{j \in C} b_{k,j}$ and $w(\lambda_k) = \sum_{j \in C} w_j(\lambda_k)$, the total demand $d(k, \lambda_k)$ is $d(k, \lambda_k) = \sum_{j \in C} d_j(k, \lambda_k) = b_k + w(\lambda_k)$. We make the following assumptions:

- The baseline demand $b_{k,j}$ and the total baseline demand b_k are exogenous, bounded, dependent on time, but independent of λ_k. For instance, for a household, the baseline demand can characterize the minimum necessary power usage, such as cooking and a minimum level of illumination.
- We assume that $w(\lambda_k)$ is a decreasing function of λ_k, which is consistent with intuition. Human-induced demand response to price change has been observed in previous studies [Allcott 2009; Sweeny 2002]. In particular, an experiment conducted with 693 household consumers [Allcott 2009] showed that households exhibited significant price elasticity, and energy management and information technology could significantly increase this elasticity. With the increasing adoption of smart appliances and home automation systems, this demand response will become more automated. For instance, in the Load Guard Automatic Price Response Service provided by ComEd [ComEd 2014], a cyber-enabled controller automatically regulates the duty cycling of a central air conditioner based on ComEd's real-time prices.
- As there are a large number of consumers, we assume that $b_{k,j}$ and $w_j(\cdot)$ are unknown to the ISO. However, the ISO knows the historical total demand $\{d(h, \lambda_h) | \forall h < k\}$.

A subset of analytic results in this paper require that the price-responsive demand model satisfies the following property:

Definition 3.1. The first derivative of the price-responsive demand model, i.e., $\dot{w}(x)$, is said to be *decomposable*, if $\dot{w}(x)|_{x=\gamma\lambda} = \dot{w}(x)|_{x=\lambda} \cdot \mu(\gamma|\Theta)$, where Θ is the set of model parameters of $w(x)$, γ and $\mu(\gamma|\Theta)$ are independent of λ. For simplicity of exposition, we denote $\mu(\gamma|\Theta)$ as μ in the rest of this paper.

3.3. Suppliers

Let S denote the set of suppliers in the system. Let $s_i(\lambda_k)$ denote the quantity of power that a supplier $i \in S$ schedules to generate in the kth pricing period given price λ_k. Let $s(\lambda_k)$ denote the *scheduled* total supply in the kth pricing period, i.e., $s(\lambda_k) = \sum_{i \in S} s_i(\lambda_k)$. We make the following assumptions:

- We assume that $s(\lambda_k)$ is an increasing function of λ_k. This assumption can be validated using published electricity market data. For instance, using half-hourly supply data of New South Wales (NSW), Australia, provided by the Australian Energy Market Operator (AEMO) [Austrilian Energy Market Operator 2014], Fig. 1 shows a histogram of the total supply versus the wholesale price in January, 2012. We can see that the total supply increases with the price. Such a monotonic relationship can also been seen in the electricity market of California [Sweeny 2002, p. 112]. Note that, in current electricity wholesale markets, the supply and price are often determined through a bidding process [Fleten and Pettersen 2005], which is generally governed

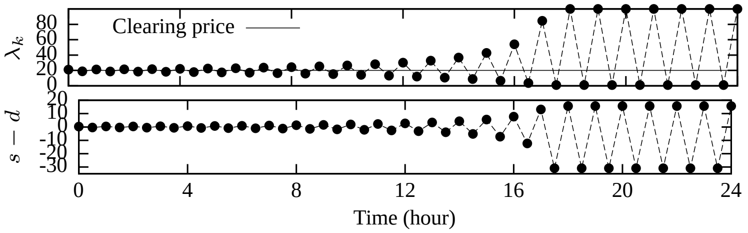

Fig. 2. An example of unstable solution [Roozbehani et al. 2012a]. Top figure: Evolution of price. Bottom figure: Generation scheduling error in GW. (Settings: $s(\lambda_k) = 152\lambda_k + 4503$, $w(\lambda_k) = 33447 \cdot \lambda_k^{-0.6}$, $b_k = 2000$, $\lambda_k^* = 20$, $\lambda_0 = 21$, $\lambda_{\min} = 1$, $\lambda_{\max} = 100$)

by generation costs. In a competitive bidding-based wholesale market, the resultant supply and price will well reflect the supply model $s(\lambda_k)$ derived from the generation cost model.

— In this paper, we consider centralized bulk generation rather than distributed generation. As there are typically a limited number of suppliers, we assume that the ISO can estimate the total supply model $s(\cdot)$.

— We assume that the generation capacity of the supplier i is at least $s_i(\lambda_{\max})$.

4. THE RTP PROBLEM AND SOLUTIONS

This section formally states the RTP problem, examines an existing solution, and describes a basic control-theoretic solution with provable *bounded-input bounded-output stability* (referred to as *stability* for short in this paper). Based on the solution, the security analysis in Sections 5 and 6 lays the foundation for understanding the impact of attacks on feedback-based RTP systems.

4.1. The RTP Problem and Solution Stability

At time instant $k{\cdot}T$, the ISO aims to find the *clearing price* for the period $[k{\cdot}T, (k{+}1){\cdot}T]$, denoted by λ_k^*, such that the scheduled supply matches demand, i.e., $s(\lambda_k^*) = d(k, \lambda_k^*)$. Existing studies often assume a known or a learned demand model (i.e., $d(k, \lambda_k)$) such that the clearing price can be solved in closed form or obtained by a search algorithm [Yu et al. 2012; Samadi et al. 2014; Choi et al. 1998; Siddiqi and Baughman 1993]. The effectiveness of these open-loop approaches rely on an accurate demand model, which is often difficult to obtain in practice. Recent studies attempt to reduce the reliance on an accurate demand model by multiple rounds of communications between the consumers and suppliers to converge to the clearing price in each pricing period [Samadi et al. 2010; Qian et al. 2013; Li et al. 2011]. However, this requires the two-way communication capability of consumers, which is still largely unavailable for end customers in today's power grids.

Researchers also start to examine the applicability of the current RTP schemes employed in today's wholesale markets to the scenario that the wholesale real-time prices are relayed to end consumers [Roozbehani et al. 2012a]. In the current practice, the ISO sets the price to match the scheduled supply and *predicted* demand (denoted by $\widetilde{d}(k, \lambda_k)$). Formally, the ISO solves the following problem:

RTP Problem: Find λ_k such that $s(\lambda_k) = \widetilde{d}(k, \lambda_k)$.

Current demand prediction algorithms often include an autoregressive component. As shown in [Roozbehani et al. 2012a], such an autoregression-based prediction may lead to significantly fluctuating prices. For instance, the simplest form of the autoregression, which uses $d(k-1, \lambda_{k-1})$ as the predicted demand $\widetilde{d}(k, \lambda_k)$, can yield oscillating prices as shown in Fig. 2. This simplest RTP scheme is called *direct feedback* approach [Roozbehani et al. 2012a] in the rest of this paper. The root cause of the oscillation is the unstable closed-loop system formed by the direct feedback.

As a fundamental requirement for any physical or economic component of a power grid, stability ensures that the component is resilient to certain exogenous disturbances and inaccuracies in the system models. These disturbances and inaccuracies are often inevitable. Therefore, instability is highly undesirable in the design and operation of a system. Particularly, in an unstable RTP system, the price set by the ISO will oscillate or diverge, even if the initial price is very close to the true clearing price, where the tiny error may be caused by exogenous factors such as temporal variation of baseline demand. The oscillations may lead to undesirable consequences. For instance, when the diverging prices reach low values, the increased demand may cause overload of the transmission and distribution networks. Moreover, as shown in Fig. 2, the unstable system may experience significant *generation scheduling errors* (i.e., $s(\lambda_k) - d(k, \lambda_k)$). If the suppliers are responsible for handling these errors, reserve generating capacities can help compensate for the errors. However, their use may increase the cost of operating the system. We note that the exact impact of the instability on the grid depends on many operational regulations of the grid, such as the approach to handling generation scheduling errors, the ISO's ability to detect price oscillation, and its follow-up mitigation policies. However, the system complexity and the increased operating cost caused by these countermeasures will offset or even invalidate the RTP's promise of improving system efficiency. Therefore, ensuring stability of RTP is a first priority for system designers and ISOs. Moreover, the fundamental conditions for ensuring the RTP stability in the presence of certain cybersecurity attacks are also of great interest, due to the increasing exposure of grids to cyberspace.

To study the impact of integrity attacks on the RTP systems, we should start with RTP schemes that are stable in the absence of attacks. Our analysis and extensive numerical experiments show that, the direct feedback approach [Roozbehani et al. 2012a] is unstable and diverging with significant probability. Moreover, according to [Roozbehani et al. 2012a], if the direct feedback approach is not stable, it is difficult to stabilize those systems based on autoregressive demand prediction. The details of the analysis and numerical experiments are omitted due to space constraints and can be found in Appendix B of the supplementary file of this paper.

4.2. Control-Theoretic Price Stabilization

The results in Section 4.1 show the necessity of control laws for stabilizing RTP systems, which was also pointed out in [Roozbehani et al. 2012a]. We note that the design of RTP algorithms to meet various specific requirements (e.g., stability, efficiency, etc) under the scenario that the end consumers are provided real-time wholesale prices is still under active research. In this paper, we do not aim to design specific RTP algorithms. Instead, we aim to analyze the impacts of integrity attacks against the vulnerable real-time price signals on the stability of the RTP systems that are stable in the absence of attacks. To achieve this, in this section, we describe a basic control-theoretic RTP algorithm as a baseline system for the security analysis in the following sections. We expect our security analysis can provide a baseline for understanding the security properties of other sophisticated RTP algorithms that are expected to have better stability. Moreover, we also hope our results will inform RTP designers to take potential security threats into consideration.

The objective of price stabilization is to minimize the generation scheduling error and adapt to the time-varying baseline load. We reformulate the RTP problem as a classical discrete-time feedback control problem. Under this formulation, the ISO observes the generation scheduling error in the previous pricing period, and then uses it

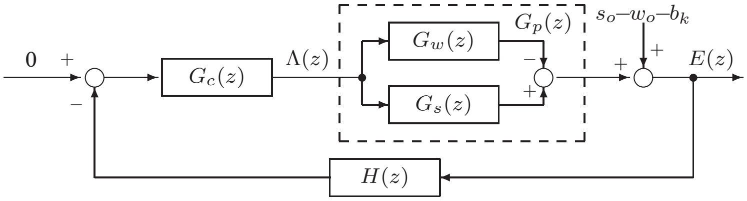

Fig. 3. The control-theoretic price stabilization. $\Lambda(z)$ and $E(z)$ are the z-transforms of λ_k and e_k.

to guide the setting of the price in the next pricing period. Specifically, let e_k denote the generation scheduling error, i.e., $e_k = s(\lambda_k) - d(k, \lambda_k)$. The objective is to maintain the *controlled variable* e_k close to its *reference*, which is zero. The *manipulated variable* is λ_k, and $s(\lambda_k) - d(k, \lambda_k)$ is the *controlled system*. The block diagram of the feedback control loop is shown in Fig. 3. We let $G_c(z)$, $G_p(z)$, and $H(z)$ denote the transfer functions of the RTP algorithm, the controlled system, and the observation system, which are expressed in the z-transform domain. The z-transform [Ogata 1995] provides a compact representation for discrete-time functions, where z represents a time shift operation. As b_k is bounded and independent of λ_k, it can be modeled as a disturbance.

We now derive the expressions of $G_p(z)$ and $H(z)$. Note that the supply and price-responsive demand models, i.e., $s(\lambda)$ and $w(\lambda)$, can be non-linear. In controller design, a common approach to dealing with non-linear systems is to adopt local linearization [Ogata 1995]. Specifically, $s(\lambda) \simeq s(\lambda_o) + \dot{s}(\lambda_o) \cdot (\lambda - \lambda_o)$ and $w(\lambda) \simeq w(\lambda_o) + \dot{w}(\lambda_o) \cdot (\lambda - \lambda_o)$, where λ_o is a *fixed operating point*. By denoting $s_p(\lambda) = \dot{s}(\lambda_o)\lambda$, $s_o = s(\lambda_o) - \dot{s}(\lambda_o)\lambda_o$, $w_p(\lambda) = \dot{w}(\lambda_o)\lambda$, and $w_o = w(\lambda_o) - \dot{w}(\lambda_o)\lambda_o$, we have $s(\lambda) \simeq s_p(\lambda) + s_o$ and $w(\lambda) \simeq w_p(\lambda) + w_o$. As s_o and w_o are independent of λ, as shown in Fig. 3, we can collect them with the price-independent b_k. The transfer functions of the proportional models $s_p(\lambda)$ and $w_p(\lambda)$ are $G_s(z) = \dot{s}(\lambda_o)$ and $G_w(z) = \dot{w}(\lambda_o)$, respectively. Therefore, $G_p(z) = G_s(z) - G_w(z) = \dot{s}(\lambda_o) - \dot{w}(\lambda_o)$. As the RTP algorithm uses the observed generation scheduling error in the previous pricing period to adjust the price for the current pricing period, $H(z) = z^{-1}$, which represents the delay of one pricing period.

The analysis in the rest of this paper is based on the above linearized abstract supply and demand models at a fixed operating point. Moreover, in Section 5, we will develop an approximate formula in Eq. (1) to characterize the total demand when a fraction of the consumers receive compromised price signals. The linearization and approximation enable us to study the RTP's susceptibility to malicious attacks under LTI settings. For systems that involve highly non-linear and/or time-variant components, the LTI treatments in this paper may lead to inaccuracies in characterizing such systems. However, the analysis in this paper provides insights into understanding the problem under real-world settings that are often non-linear; it also gives a basis for finer analysis based on piecewise linearization, as we discuss for future work in Section 9. We also note that the simulations in Section 8 are driven by a non-linear demand model introduced in Section 5.4. The simulation results are consistent with the analysis based on linearized models.

Based on the above modeling, we have the following proposition. The proof is omitted due to space constraints and can be found in Appendix E of the supplementary file of this paper.

PROPOSITION 4.1. *For the linearized system* $G_p(z) = \dot{s}(\lambda_o) - \dot{w}(\lambda_o)$ *with* λ_o *fixed and the observation system* $H(z) = z^{-1}$*, the following RTP algorithm ensures stability:* $\lambda_k = \lambda_{k-1} - \frac{2\eta}{\dot{s}(\lambda_o) - \dot{w}(\lambda_o)} \cdot e_{k-1}$*, where* $\eta \in (0, 1)$.

The intuition of the algorithm in Proposition 4.1 is as follows. Due to the monotonicity of $s(\cdot)$ and $w(\cdot)$, the coefficient $\frac{2\eta}{\dot{s}(\lambda_o)-\dot{w}(\lambda_o)}$ in the proposition is a positive constant. If the observed generation scheduling error e_{k-1} is positive (i.e., the scheduled generation is larger than the demand), the ISO should decrease the price such that the scheduled generation will decrease and the demand will increase; otherwise, the ISO should increase the price. In Proposition 4.1, the decrease or increase in price is proportional to the error e_{k-1}. This ensures that the system can converge to an equilibrium where the generation scheduling error is zero, if b_k is a constant over time. When b_k is a variable over time, the algorithm can adapt to the change of b_k. To illustrate, assume that the system has converged in the $(k-2)$th pricing period (i.e., $e_{k-2}=0$) and b_{k-1} is smaller than b_{k-2}. Then, e_{k-1} is positive due to the decreased baseline demand. As a result, the algorithm will decrease the price to in turn decrease the scheduled supply and achieve a new equilibrium. As the algorithm is bounded-input bounded-output stable, if b_k is a bounded variable over time, the e_k is also bounded.

We note that an RTP algorithm proposed by Dalkilic et al. [Dalkilic et al. 2013] also tunes the price based on the immediate past generation scheduling error. Compared with other RTP algorithms [Yu et al. 2012; Samadi et al. 2014; Choi et al. 1998; Siddiqi and Baughman 1993; Samadi et al. 2010; Qian et al. 2013; Li et al. 2011], our and Dalkilic's algorithms require neither an accurate demand model nor the two-way communication capability of consumers. The difference between the two algorithms is, in addition to ensuring stability, Dalkilic's algorithm also minimizes generation cost by a different approach to determine the gain for the generation scheduling error. As this paper focuses on the impact of attacks on the stability of RTP only, our algorithm in Proposition 4.1 is a good baseline.

From control theory, when b_k is a constant, the system converges fastest when $\eta = 0.5$, as the system's pole is at the origin. The convergence speed is important for adapting to fast time-varying baseline load so that the convergence is achieved before a significant change of baseline load. However, our analysis in Sections 5 and 6 shows that we generally need to set a smaller η to reduce the impact of attacks. In other words, we have to sacrifice convergence speed for resilience to attacks.

As discussed in Section 3.2, $w(\cdot)$ is unknown to the ISO. In practice, the ISO can estimate $\dot{w}(\lambda_o)$ based on the history of price-demand pairs. Our analysis shows that, if the relative error in estimating $\dot{w}(\lambda_o)$ is less than $100 \times (1-\eta)\%$, the algorithm given by Proposition 4.1 remains stable. The details of the analysis are omitted due to space constraints and can be found in Appendix C of the supplementary file of this paper. For instance, if $\eta = 0.5$, the relative error bound is 50%, which is a tractable requirement for most estimation algorithms. Moreover, for a smaller η that is set to increase resilience to attacks, the error bound will be larger. As the focus of this paper is to analyze the fundamental impact of integrity attacks on system stability under the control law in Proposition 4.1, we do not elaborate on the estimation algorithm, and the security analysis in Sections 5 and 6 assumes that the ISO can estimate $\dot{w}(\lambda_o)$ accurately.

5. INTEGRITY ATTACKS ON REAL-TIME PRICES FOR CONSUMERS

This section studies the impact of scaling and delay attacks on the real-time prices for consumers under the RTP scheme given by Proposition 4.1. Section 5.1 defines the attack impact metrics. Sections 5.2 and 5.3 present the analytic results for the impact of scaling and delay attacks, respectively. Section 5.4 presents numerical results obtained under specific demand and supply models. The proofs for all the propositions in this section are omitted due to space constraints and can be found in Appendix F of the supplementary file of this paper. Note that the analysis framework in this section can be extended to address the attacks on the prices for suppliers (cf. Appendix G of the supplementary file).

5.1. Attack Models and Impact Metrics

We consider integrity attacks on the price signals received by a subset of consumers. If the price signal received by a consumer is subject to attack, the price signal applied for the current pricing period (denoted by λ'_k) is different from the true price λ_k. The integrity attacks on the price signals can be launched in different ways. For instance, once the adversary has compromised the intermediate nodes in the communication network of the power grid (e.g., routers) and obtained the decryption/encryption keys held by the ISO and/or smart meters, the adversary can intercept and forge price data packets. Moreover, recent reverse engineering and penetration tests [McLaughlin et al. 2010; Rouf et al. 2012] have shown that many smart meters lack basic security measures to ensure integrity and authenticity of the input/output data. These security vulnerabilities can be exploited to maliciously change the price signals. We would like to point out that the integrity attacks do pose strong requirements for the adversary. They require that the adversary is able to modify the price information, either at the source, during transmissions, or at the smart meters. However, these attacks in a cyber environment are certainly feasible and credible, and it would be wrongfully complacent to ignore their possibility.

5.1.1. Attack Models. As the number of consumers in a grid is often large, the number of compromised consumers is an important metric for the adversary's capability and resource availability. Let C' denote the set of consumers whose price signals are compromised, where $C' \subseteq C$, and $w'(\lambda_k)$ denote the total price-responsive demand in the presence of an attack. Thus, $w'(\lambda_k) = \sum_{j \in C'} w_j(\lambda'_k) + \sum_{j \in C \setminus C'} w_j(\lambda_k)$. We define $\rho = \frac{\sum_{j \in C'} w_j(\lambda'_k)}{\sum_{j \in C} w_j(\lambda'_k)} = \frac{\sum_{j \in C'} w_j(\lambda'_k)}{w(\lambda'_k)}$, which characterizes the fraction of consumers receiving the compromised price signals. If the consumers are homogeneous (i.e., $w_j(\cdot)$ is same for all j), ρ is a constant, i.e., $\rho = |C'|/|C|$. If they are heterogeneous, ρ is a function of λ'_k. The extensive numerical evaluation in Appendix D of the supplementary file of this paper shows that $\rho \simeq |C'|/|C|$ with a variation of less than 0.003 and hence, it can be practically treated as a constant. Moreover, we make the following approximation: $\sum_{j \in C \setminus C'} w_j(\lambda_k) \simeq (1-\rho) \sum_{j \in C} w_j(\lambda_k) = (1-\rho) w(\lambda_k)$. The numerical evaluation in Appendix D of the supplementary file shows that the relative approximation error of the above approximation is less than 1%. Therefore, we have

$$w'(\lambda_k) \simeq \rho w(\lambda'_k) + (1-\rho) w(\lambda_k). \tag{1}$$

If the price signals can be arbitrarily modified, the capability requirements of an adversary would be high. In this paper, we consider "constrained" integrity attacks, where the malicious modifications follow certain rules and can be realized with lower capability and resource requirements. Note that the adversary must be able to cause more severe damage to the system if she is assumed to be able to modify the price signals arbitrarily. An attack can be characterized by the parameters for the rule, which is denoted by $\mathcal{A}$. We consider two kinds of integrity attacks.

Definition 5.1. Under the *scaling attack* $\mathcal{A} = (\rho, \gamma)$, the compromised price is a scaled version of the true price, i.e., $\lambda'_k = \gamma \lambda_k$, $\gamma \in \mathbb{R}^+$.

Definition 5.2. Under the *delay attack* $\mathcal{A} = (\rho, \tau)$, the compromised price is an old price, i.e., $\lambda'_k = \lambda_{k-\tau}$, $\tau \in \mathbb{Z}^+$.

These two attacks can be launched in various ways. The price values or time stamps in data packets sent to the smart meters can be maliciously modified during transmissions in vulnerable communication networks. Moreover, they can be launched in indirect ways. For instance, the delay attack can be launched by modifying the smart meters' internal clocks. Smart meters typically assign a memory buffer to store received prices. If a smart meter's clock has a lag, it will likely store newly received prices in the buffer and apply an old price for the present. Furthermore, attacks on the clocks can be realized by compromising vulnerable time synchronization protocols or the time servers that provide timing information to the smart meters. A few smart meter products [Schneider Electric 2014] synchronize their clocks via a built-in GPS receiver, which is vulnerable and subject to remote attacks that are effective across large geographic areas [Nighswander et al. 2012].

In this paper, we assume that at most one kind of attack is in effect. Moreover, we assume that the attack parameters are the same for all the compromised consumers. For instance, if a delay attack with $\tau = 2$ is launched, all the compromised consumers experience the same delay of two pricing periods. These simplifications allow us to better understand the impact of each attack on the RTP system, which is the basis for understanding more complex scenarios such as heterogeneous attack parameters and combinations of attack types. In Section 7.2, we will briefly discuss how to extend our analysis to address these more complex cases.

5.1.2. Attack Impact Metrics. This section defines two metrics for the impact of the integrity attacks on system stability. We first define the *marginal demand-supply ratio*, which is a quantity that can significantly affect the system stability under attacks.

Definition 5.3. Marginal demand-supply ratio is $h = \left|\frac{\dot{w}(\lambda_o)}{\dot{s}(\lambda_o)}\right|$.

From Definition 5.3, h depends on the operating point λ_o. As discussed in Section 4.2, the gain coefficient η of the RTP algorithm affects the system stability in a major way. Therefore, we define the following metric:

Definition 5.4. Given attack $\mathcal{A}$, the *region of operating point stability* under attack, denoted by $\mathrm{ROS}_{\lambda_o}(\mathcal{A})$, is $\mathrm{ROS}_{\lambda_o}(\mathcal{A}) = \{(h, \eta) | \text{The system is stable under attack } \mathcal{A}\}$.

The above metric depends on λ_o. We define a second metric that is independent of λ_o:

Definition 5.5. Given attack $\mathcal{A}$, the *region of stability* under attack, denoted by $\mathrm{ROS}(\mathcal{A})$, is $\mathrm{ROS}(\mathcal{A}) = \{\eta | \text{The system is stable under attack } \mathcal{A}, \forall h > 0\}$.

The above two metrics are important for understanding the impact of integrity attacks on the stability of the RTP system under the RTP algorithm in Proposition 4.1. In particular, the $\mathrm{ROS}(\mathcal{A})$ specifies the range of η that ensures system stability under attack $\mathcal{A}$. Hence, the ROS allows us to compare the impacts of different integrity attacks. For two attacks $\mathcal{A}_1$ and $\mathcal{A}_2$, if $\mathrm{ROS}(\mathcal{A}_1) \subset \mathrm{ROS}(\mathcal{A}_2)$, the ISO has more flexibility in setting η under $\mathcal{A}_2$ than $\mathcal{A}_1$, to achieve faster convergence. Thus, the system is more resilient to $\mathcal{A}_2$ than $\mathcal{A}_1$. From the adversary's perspective, $\mathcal{A}_1$ is more effective than $\mathcal{A}_2$. Note that, when the RTP system with $\eta \in \mathrm{ROS}(\mathcal{A})$ is stable under attack $\mathcal{A}$, the compromised consumers may still experience monetary losses and the system may run at low efficiency. However, this paper focuses on the impact of attacks on the system stability, which is a fundamental system requirement. In Sections 5.2 and 5.3, we will derive the ROS_{λ_o} and ROS for the scaling and delay attacks.

5.2. Impact of Scaling Attacks on Prices for Consumers

The local linearization of Eq. (1) with $\lambda'_k = \gamma\lambda_k$ is

$$w'(\lambda_k) \simeq \rho \cdot (w(\gamma\lambda_o) + \dot{w}(x)|_{x=\gamma\lambda_o} \cdot (\gamma\lambda_k - \gamma\lambda_o)) + (1-\rho) \cdot (w(\lambda_o) + \dot{w}(\lambda_o) \cdot (\lambda_k - \lambda_o)).$$

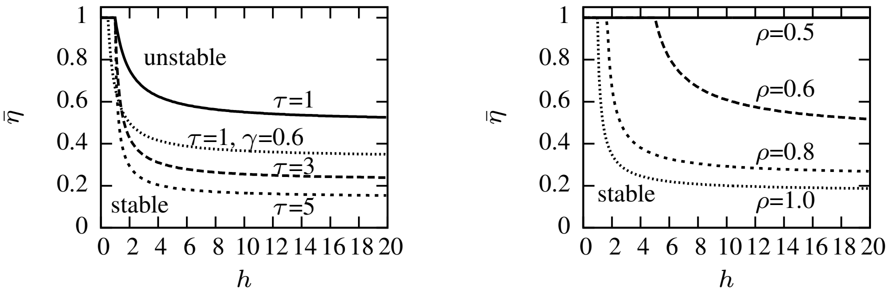

(a) $\rho = 1$

(b) $\tau = 4$

Fig. 4. Stability boundaries under delay attack on prices for consumers. (In the left figure, the curve with $\tau = 1$ and $\gamma = 0.6$ is for combined attack discussed in Section 7.2.)

By collecting the price-independent terms with b_k, the transfer function of the price-dependent component is $G_w(z) = \rho\gamma\dot{w}(x)|_{x=\gamma\lambda_o} + (1-\rho)\dot{w}(\lambda_o)$. The propositions in the following two subsections give the region of operating point stability and region of stability when $\dot{w}(\cdot)$ is decomposable as defined in Definition 3.1. The proofs can be found from Appendix F.1 and F.2 of the supplementary file of this paper.

PROPOSITION 5.6. *For the scaling attacks on the prices for consumers and the linearized system based on a fixed operating point λ_o and a decomposable $\dot{w}(\cdot)$, $\mathrm{ROS}_{\lambda_o}(\rho,\gamma) = \{(h,\eta)|0 < \eta < \min\{1,\bar{\eta}\}, \forall h > 0\}$, where $\bar{\eta} = \frac{h+1}{h+1+\rho h(\gamma\mu-1)}$ and μ is defined in Definition 3.1.*

PROPOSITION 5.7. *For the scaling attacks on the prices for consumers and the linearized system based on a decomposable $\dot{w}(\cdot)$, when $\gamma\mu \in (0,1]$, $\mathrm{ROS}(\rho,\gamma) = \{\eta|0 < \eta < 1\}$; when $\gamma\mu > 1$, $\mathrm{ROS}(\rho,\gamma) = \{\eta\,|0 < \eta < \inf_{h>0}\bar{\eta}\}$, where $\inf_{h>0}\bar{\eta} = \frac{1}{1+\rho(\gamma\mu-1)}$. Note that μ is defined in Definition 3.1.*

5.3. Impact of Delay Attacks on Prices for Consumers

The local linearization of Eq. (1) with $\lambda'_k = \lambda_{k-\tau}$ is

$$w'(\lambda_k) \simeq \rho \cdot (w(\lambda_o) + \dot{w}(\lambda_o)\cdot(\lambda_{k-\tau} - \lambda_o)) + (1-\rho)\cdot(w(\lambda_o) + \dot{w}(\lambda_o)\cdot(\lambda_k - \lambda_o)).$$

By collecting the price-independent terms with b_k, the transfer function of the price-dependent component is $G_w(z) = z^{-\tau}\rho\dot{w}(\lambda_o) + (1-\rho)\dot{w}(\lambda_o)$, where $z^{-\tau}$ represents a delay of τ pricing periods. Therefore, $G_p(z) = G_s(z) - G_w(z) = \dot{s}(\lambda_o) - z^{-\tau}\rho\dot{w}(\lambda_o) - (1-\rho)\dot{w}(\lambda_o)$. The closed-loop transfer function under the attack is $T_c(z) = \frac{G_c(z)G_p(z)}{1+G_c(z)G_p(z)H(z)} = \frac{2\eta(1+(1-\rho)h)z^{\tau+1}+2\rho\eta hz}{P(z)}$, where the system characteristic function is $P(z) = (h+1)z^{\tau+1} + (2\eta + 2\eta(1-\rho)h - h - 1)\,z^{\tau} + 2\eta\rho h$.

5.3.1. Region of Operating Point Stability. As $P(z)$ is a $(\tau+1)$-order polynomial, it is extremely difficult to derive the closed-form formulas for the poles of $T_c(z)$. Various methods have been developed to test the stability without explicitly solving for the poles [Ogata 1995]. Among them, the Jury test [Ogata 1995, p. 185] is preferred because the coefficients of $P(z)$ are real numbers. The Jury test constructs a table based on the coefficients of $P(z)$ and derives the stability conditions from the table. Given ρ, we can derive the closed-form ROS_{λ_o} for different τ from the Jury test. However, the expressions become more complicated for larger τ. We numerically compute the ROS_{λ_o} based on the Jury test for various settings of τ and ρ. Fig. 4 plots the stability boundaries under various settings of h, τ, and ρ, where the ROS_{λ_o} are the regions below the

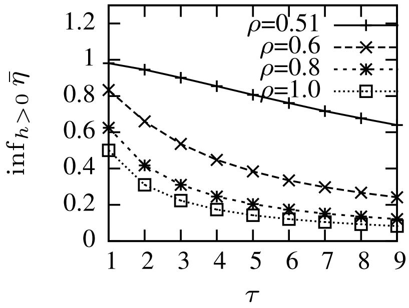

Fig. 5. Upper bound of ROS under delay attacks on prices for consumers.

boundaries. From Fig. 4, the ROS_{λ_o} shrinks with τ and ρ, which is consistent with intuition. We have the following proposition. The proof is based on the Jury test, which can be found in Appendix F.3 of the supplementary file of this paper.

PROPOSITION 5.8. *For the delay attacks on the prices for consumers and the linearized system with a fixed operating point* λ_o*,* $\mathrm{ROS}_{\lambda_o}(\rho, \tau+1) \subseteq \mathrm{ROS}_{\lambda_o}(\rho, \tau)$.

5.3.2. Region of Stability. We observe from Fig. 4(b) that, when $\rho = 0.5$, the system is stable for $\eta \in (0, 1)$. We have the following proposition.

PROPOSITION 5.9. *For the delay attacks on the prices for consumers and the linearized system, if* $\rho \in (0, 0.5]$*,* $\forall \tau \in \mathbb{Z}^+$*,* $\mathrm{ROS}(\rho, \tau) = \{\eta | 0 < \eta < 1\}$.

The proof can be found in Appendix F.4 of the supplementary file, where we prove that if $\rho \in (0, 0.5]$, all the roots of $P(z)$ are within the unit circle centered at the origin in the z-plane and hence the system is stable [Ogata 1995]. From Proposition 5.9, to launch a successful delay attack that destabilizes the system, the adversary has to compromise no less than half of the consumers. The intuition behind this result is that the compromised price-responsive load must predominate to affect the operation of the system. This result poses strong requirements for the adversary. However, she could accomplish the goal by targeting shared infrastructures such as time servers that provide timing information to all the smart meters. On the other hand, the need for the adversary to compromise a large fraction of the meters in order to be effective is indicative of the resilience of the RTP algorithm in Proposition 4.1 to delay attacks.

We now discuss the ROS when $\rho \in (0.5, 1]$. From Fig. 4, the stability boundary curves are non-increasing and converge to limits when $h \to +\infty$. Let $\bar{\eta}(h|\rho, \tau)$ denote the stability boundary curve for particular ρ and τ. Therefore, $\mathrm{ROS}(\rho, \tau) = \{\eta | 0 < \eta < \lim_{h \to +\infty} \bar{\eta}(h|\rho, \tau)\}$. When $\tau = 1$, the limit is simply $\frac{1}{2\rho}$. However, for larger τ, it is extremely difficult to derive the closed-form formula for the limit, primarily because of the iterative nature of the Jury test. We have developed a symbolic algorithm to define $\mathrm{ROS}(\rho, \tau)$ based on key observations from the Jury test procedure. The algorithm is omitted in this paper and can be found in Algorithm 1 of the supplementary file of this paper. Fig. 5 plots $\lim_{h \to +\infty} \bar{\eta}(h|\rho, \tau)$, which is computed by the algorithm, versus τ under various settings of ρ. From the figure, we can see that the ROS shrinks with ρ and τ, which is consistent with intuition.

5.4. Numerical Results based on Specific Demand and Supply Models

The previous sections presented analytic results based on the general demand and supply models. To illustrate and validate these results, this section presents numerical results obtained under specific instantiations of the demand and supply models,

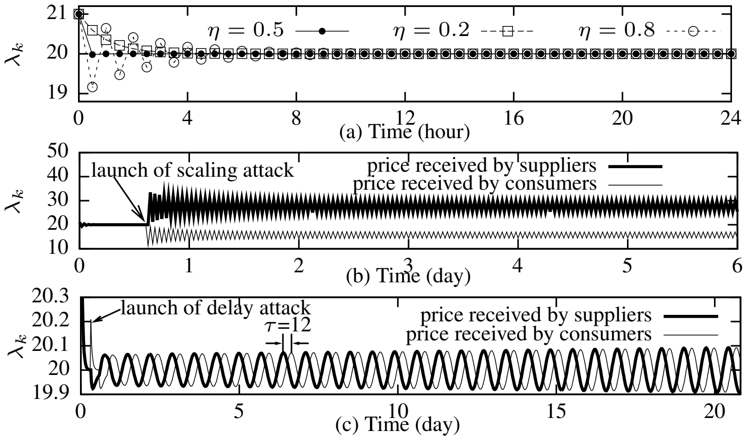

Fig. 6. A numerical example under the linear supply and CEO demand models (Settings: $T = 0.5$, p=152, q=4503, D=60893, ϵ=−0.8, b_k=2000, λ^*=20, λ_0=21). (a) Price stabilization; (b) Scaling attack on prices for consumers (η=0.8, ρ=1, γ=0.57); (c) Delay attack on prices for consumers (η=0.2, ρ=1, τ=12).

namely a constant elasticity of own-price (CEO) demand model, a linear supply model, and a quadratic supply model. We now define these three models.

— The CEO model [Fleten and Pettersen 2005] is defined by $w(\lambda_k) = D \cdot \lambda_k^{\epsilon}$, where D and ϵ are positive and negative constants, respectively. The ϵ, often referred to as *price elasticity of demand*, is typically within $(-1, 0)$ [Filippini 2011; Lijesen 2007]. In this section, we set $D = 60893$ and $\epsilon = -0.8$. Moreover, $\dot{w}(\cdot)$ is decomposable and $\mu = \gamma^{\epsilon-1}$. Note that as μ under the CEO model is independent of D, all the analytic results in Sections 5.2 and 5.3 are independent of D as well.

— The linear supply model is defined by $s(\lambda_k) = p \cdot \lambda_k + q$, where p and q are two constants. A linear fitting to the supply and price data in Fig. 1 yields $p = 152$ and $q = 4503$, which are adopted in this section.

— The quadratic supply model is defined by $s(\lambda_k) = p_2 \cdot \lambda_k^2 + p_1 \cdot \lambda_k + p_0$, where p_2, p_1, and p_0 are three constants. A quadratic curve fitting to the data in Fig. 1 yields $p_2 = -3.0478$, $p_1 = 329.89$, and $p_0 = 2095.2$, which are adopted in this section.

5.4.1. CEO Demand Model and Linear Supply Model. We first discuss the numerical results obtained under the CEO demand and linear supply models.

Numerical Example of Price Stabilization. We start by a numerical example to illustrate the control-theoretic price stabilization algorithm in Proposition 4.1. The algorithm assumes a fixed operating point λ_o. However, intuitively, if the operating point λ_o adapts to the current price, the linear approximations to $s(\lambda)$ and $w(\lambda)$ are more accurate. Specifically, by setting $\lambda_o = \lambda_{k-1}$, we have the following algorithm:

$$\lambda_k = \lambda_{k-1} - \frac{2\eta}{\dot{s}(\lambda_{k-1}) - \dot{w}(\lambda_{k-1})} \cdot e_{k-1}. \tag{2}$$

Although there is a lack of rigorous theory to support the technique of adapting λ_o to the current price, our numerical experiments show that the algorithm in Eq. (2) is always stable under all model parameter settings for evaluating the direct feedback approach in Section 4.1. The numerical examples and simulations conducted in the rest of this paper employ the algorithm in Eq. (2). Fig. 6(a) shows the evolution of price with fixed baseline load. When $\eta = 0.5$, λ_k converges to λ^* after two pricing periods. When $\eta = 0.2$, the system has a longer settling time. When $\eta = 0.8$, the price oscillates

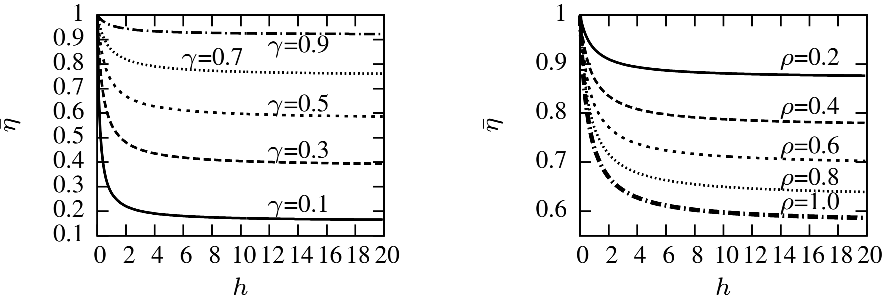

(a) $\rho = 1$

(b) $\gamma = 0.5$

Fig. 7. Stability boundaries under scaling attack on prices for consumers and CEO demand model ($\epsilon = -0.8$, ROS_{λ_o} are the regions below the boundaries).

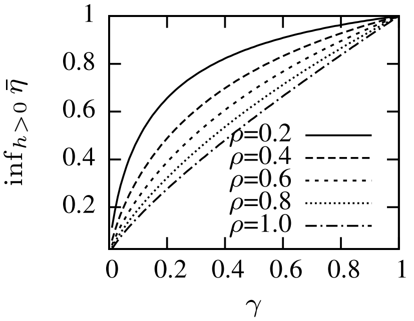

Fig. 8. Upper bound of ROS under scaling attacks on prices for consumers and CEO demand model ($\epsilon = -0.8$).

but converges. The oscillation is caused by a negative pole. Fig. 6 will also be used as a running example in the rest of this section to illustrate the impact of attacks.

Numerical Results of Scaling Attacks. We first present the numerical results of ROS_{λ_o} for scaling attacks. By replacing $\mu = \gamma^{\epsilon-1}$ in Proposition 5.6, we have $\bar{\eta} = \frac{h+1}{h+1+\rho h(\gamma^{\epsilon}-1)}$. Fig. 7 plots the stability boundaries, where the ROS_{λ_o} are the regions below the boundaries. We can see that the ROS_{λ_o} shrinks with increased ρ and decreased γ. This can be easily proved by the monotonicity of $\bar{\eta}$. Moreover, it is consistent with the intuitions that (i) the system becomes more unstable when more consumers are compromised, and (ii) the increased demand due to a decreased γ poses more challenges to the system.

We also use the numerical example in Fig. 6(b) to verify our analysis. Fig. 6(b) shows the price signals received by the suppliers and consumers, respectively, when $\gamma = 0.57$. We can see that the price does not converge. The average value of h is 0.850, which falls in the unstable region ($h > 0.786$) according to the analytical ROS_{λ_o}. Note that when $\gamma = 0.59$, the price converges and the average value of h is 0.862, which falls in the stable region ($h < 0.908$) according to the numerical results of ROS_{λ_o} in Fig. 7. Therefore, Proposition 5.6 successfully characterizes the critical stability boundary. Note that, as the settings for Fig. 6(b) are close to the stability boundary, the price oscillates in a small range. For smaller γ, the price can severely oscillate, as shown in Section 8.

We then present the numerical results of ROS for scaling attacks. Replacing $\mu = \gamma^{\epsilon-1}$ in Proposition 5.7 yields the following result. When $\gamma \geq 1$, $\mathrm{ROS}(\rho, \gamma) = \{\eta | 0 <$

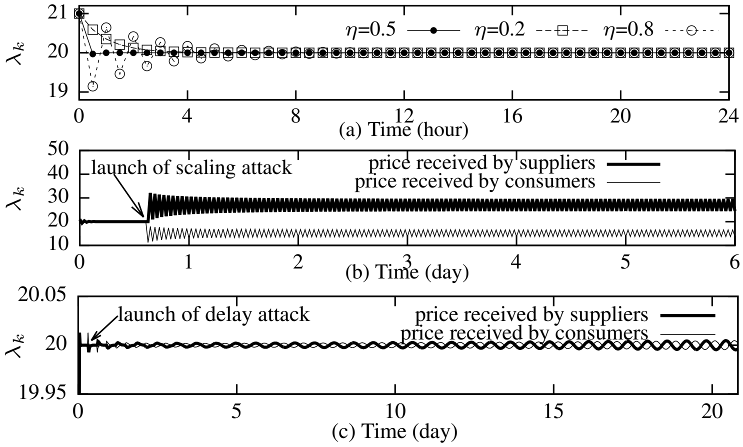

Fig. 9. A numerical example under the quadratic supply and CEO demand models (Settings: T=0.5, p_2=−3.0478, p_1=329.89, p_0=2095.2, ϵ=−0.8, b_k=2000, λ^*=20, λ_0=21). (a) Price stabilization; (b) Scaling attack on prices (η=0.8, ρ=1, γ=0.56); (c) Delay attack on prices (η=0.55, ρ=1, τ=13).

$\eta < 1\}$; when $\gamma \in (0,1)$, $\mathrm{ROS}(\rho,\gamma) = \{\eta \,|0 < \eta < \inf_{h>0} \bar{\eta}\}$, where $\inf_{h>0} \bar{\eta} = \frac{1}{1+\rho(\gamma^{\epsilon}-1)}$. Therefore, if the adversary amplifies the price, the system remains stable. This result is consistent with the intuition that decreased demand due to the amplified price poses no challenges to the system. Fig. 8 plots $\inf_{h>0} \bar{\eta}$. We can see that ROS shrinks with increased ρ and decreased γ. This can be proved by the monotonicity of $\inf_{h>0} \bar{\eta}$.

Numerical Example of Delay Attacks. We use the numerical example in Fig. 6 to verify our analysis in Section 5.3. Fig. 6(c) shows the price signals received by the suppliers and consumers, respectively, when $\eta = 0.2$, $\rho = 1$, $\tau = 12$. We can see that the price diverges. The average value of h is 1.455390, which falls in the unstable region ($h > 1.447$) according to the Jury test approach presented in Section 5.3.1. Note that when $\tau = 11$, the price converges and the average value of h is 1.455335, which falls in the stable region ($h < 1.522$) according to the Jury test. Thus, the Jury test successfully characterizes the critical stability boundary. As the settings for Fig. 6(c) are close to the stability boundary, the price diverges slowly. For larger τ, the price can diverge quickly.

5.4.2. CEO Demand Model and Quadratic Supply Model. This section presents the numerical results obtained under the CEO demand the quadratic supply models. Fig. 9(a) shows evolution of price in the absence of attacks, which is similar to Fig. 6(a).

Fig. 9(b) shows the price signals received by the suppliers and consumers, respectively, when $\gamma = 0.56$. We can see that the price does not converge. The average value of h is 0.792, which falls in the unstable region ($h > 0.734$) according to the analytical ROS_{λ_o}. Note that when $\gamma = 0.58$, the price converges and the average value of h is 0.789, which falls in the stable region ($h < 0.844$) according to the numerical results of ROS_{λ_o} in Fig. 7.

Fig. 6(c) shows the price signals received by the suppliers and consumers, respectively, when $\eta = 0.55$, $\rho = 1$, $\tau = 13$. We can see that the price diverges. The average value of h is 1.053, which falls in the unstable region ($h > 1.035$) according to the Jury test approach presented in Section 5.3.1. Note that when $\tau = 11$, the price converges and the average value of h is 1.053, which falls in the stable region ($h < 1.062$) according to the Jury test.

In summary, under the scaling (or delay) attacks, the average values of h under critical settings of γ (or τ) that lead to marginally converging and diverging prices, fall

Table I. Mean absolute percentage error (MAPE) of the linear regression model.

R_S	0			1			2			3		
R_λ	0	1	2	0	1	2	0	1	2	0	1	2
MAPE (%)	12.61	12.60	12.60	2.05	1.97	1.92	1.06	1.02	1.02	1.06	1.02	1.02

in the stable and unstable regions given by the analytic ROS_{λ_o}. Hence, Proposition 5.6 and the Jury test accurately characterize the critical stability boundaries.

6. ADDRESSING STATEFUL SUPPLIERS

Previous sections are based on a *stateless* supply model (i.e., $s(\lambda_k)$), in which the scheduled generation depends on the price only. This model captures the general supply-versus-price trend, as shown in Fig. 1. However, in practice, the generation has various limiting factors such as ramp-up constraints. Their effect can also be observed from Fig. 1. Specifically, in the data set used for plotting the figure, with a certain price, the total supply in different pricing periods can be different, leading to considerable variances as represented by the error bars in the figure. In this section, we aim to investigate the impact of these limiting factors in generation on the analytic results obtained in previous sections. However, the attempt to model each possible limiting factor would be tedious and even intractable. In this section, we will learn their aggregated effect, through an extensive empirical study to the supply-versus-price data from AEMO [Austrilian Energy Market Operator 2014] for one year. The study yields a new *stateful* supply model, which can well capture the variances observed in Fig. 1. Then, we extend our analysis to address this new supply model. The proofs for all the propositions in this section are omitted due to space constraints and can be found in Appendix H of the supplementary file of this paper.

6.1. Stateful Supply Model

As discussed in Section 5.4, the linear supply model can well approximate the relationship between the total supply and price. Motivated by this, in the following empirical study, we employ a linear regression model for total supply with an autoregressive component. Specifically, $\mathcal{S}_k = \alpha_0 + \sum_{i=1}^{R_S} \alpha_i \mathcal{S}_{k-i} + \sum_{i=0}^{R_\lambda} \beta_i \lambda_{k-i}$, where $\mathcal{S}_k$ is the total supply in the kth pricing period, R_S and R_λ are non-negative integers, $\{\alpha_i | i \in [0, R_S]\}$ and $\{\beta_i | i \in [0, R_\lambda]\}$ are the coefficients to be fitted. We conduct extensive regression analysis with different settings of R_S and R_λ based on a data set consisting of half-hourly supply-versus-price pairs from AEMO [Austrilian Energy Market Operator 2014] in the whole year of 2012. The mean absolute percentage errors (MAPE) of the fitted models are shown in Table I. Note that with the settings $R_S = 0$ and $R_\lambda = 0$, the regression model reduces to the stateless linear supply model in Section 5.4. From Table I, we can see that, including the immediate previous supply into the model (i.e., by setting $R_S = 1$) can improve the model accuracy significantly. However, including history prices (i.e., by setting $R_\lambda \geq 1$) just improves the model accuracy slightly. Thus, the regression model with the settings $R_S = 1$ and $R_\lambda = 0$ achieves a satisfactory trade-off between model complexity and accuracy. Based on this observation, in this section, we adopt a general stateful supply model as follows:

$$\mathcal{S}_k = \alpha \cdot \mathcal{S}_{k-1} + (1 - \alpha) \cdot s(\lambda_k), \tag{3}$$

where $s(\lambda_k)$ is the stateless supply model in Section 3.3, and $\alpha \in [0, 1)$ is a constant. In Eq. (3), the weights (i.e., α and $(1 - \alpha)$), which sum to one, ensure that the model preserves the trend captured by the stateless model $s(\cdot)$.

6.2. Impact of Memory Effect in Supply on the Analytic Results in Section 5

Although we can develop a new pricing algorithm based on the stateful supply model to replace the one presented in Proposition 4.1, the resultant algorithm will require knowledge about α. In this section, we are interested in how the memory effect in the stateful supply model affects our previous analysis. Specifically, we will investigate, i) whether the RTP algorithm in Proposition 4.1 developed without the consideration of supply memory effect can tolerate the effect in the absence of attacks, and ii) whether our main results on the impact of attacks (especially, Propositions 5.7 and 5.9) still hold. Under the stateful supply model in Eq. (3), the approach to handling generation scheduling errors will affect the analysis. As two extreme cases, if the consumers or suppliers fully compensate the errors, the scheduled total supply $\mathcal{S}_k$ will be $\alpha \cdot \mathcal{S}_{k-1} + (1-\alpha) \cdot s(\lambda_k)$ or $\alpha \cdot d_{k-1} + (1-\alpha) \cdot s(\lambda_k)$, respectively. We note that, the power grids nowadays often adopt a *supply-follows-demand* scheme, in which, the suppliers compensate the scheduling errors. In the grids of tomorrow, the consumers may be responsible for handling a portion of the scheduling error through demand response programs such as load curtailment. Sections 6.2.1 and 6.2.2 study these two extreme cases. The results provide insights into understanding the cases where the consumers and suppliers jointly compensate the errors. Section 6.2.3 summarizes our results in this section and discusses memory effect in demand.

6.2.1. Case 1: Consumers Fully Handle Generation Scheduling Errors.

(1) Tolerance of RTP algorithm. By following the linearization approach in Section 4.2, the stateful model in Eq. (3) can be approximated by $\mathcal{S}_k \simeq \alpha \mathcal{S}_{k-1} + (1-\alpha)\dot{s}(\lambda_o)\lambda_k - (1-\alpha)\dot{s}(\lambda_o)\lambda_o$ and the corresponding transfer function is $G_s(z) = (1-\alpha)\dot{s}(\lambda_o)/(1-\alpha z^{-1})$. Therefore, in the absence of attack, $G_p(z) = G_s(z) - G_w(z) = (1-\alpha)\dot{s}(\lambda_o)/(1-\alpha z^{-1}) - \dot{w}(\lambda_o)$. The following proposition shows that the pricing algorithm in Proposition 4.1 can tolerate the memory effect in supply.

PROPOSITION 6.1. *Under the linearized system $G_p(z) = (1-\alpha)\dot{s}(\lambda_o)/(1-\alpha z^{-1}) - \dot{w}(\lambda_o)$, the pricing algorithm in Proposition 4.1 ensures stability.*

(2) Impact of Scaling Attacks. We first derive the region of operating point stability under the stateful supply model. We have the following proposition.

PROPOSITION 6.2. *Under the stateful supply model in Eq. (3), Proposition 5.6 holds with a new $\bar{\eta}$ given by $\bar{\eta} = \frac{(\alpha+1)(h+1)}{(\alpha+1)(\rho(\gamma\mu-1)+1)h-\alpha+1}$.*

If $\alpha = 0$, the stateful supply model in Eq. (3) reduces to the stateless supply model $s(\lambda_k)$, and Proposition 6.2 reduces to Proposition 5.6.

By simply following the procedure in the proof of Proposition 5.7 with the new $\bar{\eta}$ in Proposition 6.2, we have the following proposition, which shows that the memory effect in supply does not change our main result for scaling attack.

PROPOSITION 6.3. *Under the stateful supply model Eq. (3), Proposition 5.7 holds.*

(3) Impact of Delay Attacks. By following the analysis approach in Section 5.3, the system characteristic function under the delay attack is given by

$$P(z) = (h+1)z^{\tau+2} + (-2\eta h\rho + 2\eta h - \alpha h - h - 2\alpha\eta + 2\eta - \alpha - 1)z^{\tau+1} + (2\alpha\eta h\rho - 2\alpha\eta h + \alpha h + \alpha)z^{\tau} + 2\eta h\rho z - 2\alpha\eta h\rho.$$

We numerically compute the ROS_{λ_o} based on the Jury test for various settings of α and ρ. Fig. 10(a) plots the stability boundaries, where the ROS_{λ_o} are the regions below the boundaries. We can see that the ROS_{λ_o} shrinks with α. Intuitively, due to the maliciously introduced delays, the price-inelasticity induced by the memory effect in supply makes the system less capable to respond to generation scheduling errors in

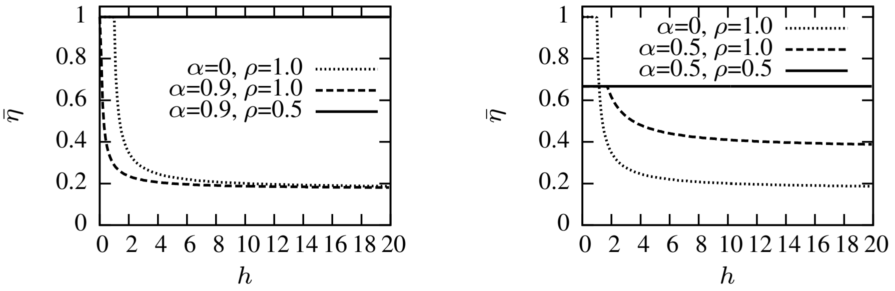

(a) Case 1: Consumers handle generation scheduling error

(b) Case 2: Suppliers handle generation scheduling error

Fig. 10. Stability boundaries under delay attack on price for consumers and stateful supply model ($\tau = 4$).

time, thus downgrading the system resilience. Moreover, we can see that when $\rho = 0.5$, the system is stable for $\eta \in (0, 1)$. However, the proof of Proposition 5.9 under the stateful supply model is difficult, because the much more complicated $P(z)$ makes it difficult to construct a function $g(A)$ (cf. proof of Proposition 5.9) that can test the magnitudes of poles. Extensive numerical experiments validate Proposition 5.9 under the stateful supply model, $\forall \alpha \in (0, 1)$ and $\forall \tau \in [1, 48]$. Moreover, under the stateful supply model, numerical experiments yield a figure same to Fig. 5, regardless of α. This suggests that the memory effect in supply does not affect the ROS in the presence of delay attacks.

6.2.2. Case 2: Suppliers Fully Handle Generation Scheduling Errors.

(1) Tolerance of RTP Algorithm. If suppliers handle the errors, the realized generation in the previous pricing period is d_{k-1}. Thus, the scheduled total supply is $\mathcal{S}_k = \alpha \cdot (b_{k-1} + w(\lambda_{k-1})) + (1 - \alpha) \cdot s(\lambda_k)$. By following the analysis approach in Section 4.2, $G_p(z)$ can be derived as $G_p(z) = \alpha \dot{w}(\lambda_o) z^{-1} + (1-\alpha)\dot{s}(\lambda_o) - \dot{w}(\lambda_o)$. We have the following proposition.

PROPOSITION 6.4. *Under the linearized system $G_p(z) = \alpha \dot{w}(\lambda_o) z^{-1} + (1 - \alpha)\dot{s}(\lambda_o) - \dot{w}(\lambda_o)$, the pricing algorithm in Proposition 4.1 ensures stability if $\eta \in (0, \frac{1}{1+\alpha})$.*

As $\frac{1}{1+\alpha} > 0.5$, a sufficient stability condition for the pricing algorithm is $\eta \in (0, 0.5]$, which can be used if α is unknown or cannot be estimated accurately. From Proposition 6.4, the memory effect in supply under Case 2 changes the stability condition of the pricing algorithm, in contrast to Case 1. Under Case 2, the memory effect creates an additional closed loop between supply and demand (i.e., $\mathcal{S}_k$ depends on d_{k-1}), which changes the system structure and thus the stability condition.

(2) Impact of Scaling Attacks. The region of operating point stability under the stateful supply model is given by the following proposition.

PROPOSITION 6.5. *Under the stateful supply model in Eq. (3), the ROS_{λ_o} given by Proposition 5.6 is revised as $\mathrm{ROS}_{\lambda_o} = \{(h, \eta) | 0 < \eta < \min\{\frac{1}{1+\alpha}, \bar{\eta}\}, \forall h > 0\}$, where $\bar{\eta} = \frac{h+1}{(\alpha+1)(\rho(\gamma\mu-1)+1)h-\alpha+1}$ and μ is defined in Definition 3.1.*

If $\alpha = 0$, Proposition 6.5 reduces to Proposition 5.6. By simply following the procedure in the proof of Proposition 5.7 with the new $\bar{\eta}$ in Proposition 6.5, we have the following proposition.

PROPOSITION 6.6. *Under the stateful supply model and the linearized system based on a decomposable $\dot{w}(\cdot)$, when $\gamma\mu \in (0,1]$, $\mathrm{ROS}(\rho,\gamma) = \{\eta | 0 < \eta < \frac{1}{1+\alpha}\}$; when $\gamma\mu > 1$, $\mathrm{ROS}(\rho,\gamma) = \{\eta | 0 < \eta < \inf_{h>0} \bar{\eta}\}$, where $\inf_{h>0} \bar{\eta} = \frac{1}{(1+\alpha)\cdot(1+\rho(\gamma\mu-1))}$. Note that μ is defined in Definition 3.1.*

Since $(1, \frac{1}{1+\alpha})$ is the new stability condition for the pricing algorithm (cf. Proposition 6.4), Proposition 6.6 shows that the memory effect in supply does not change the implication of our main result for the scaling attack discussed in Section 5.4.1.

(3) Impact of Delay Attacks. By following the analysis approach in Section 5.3, the closed-loop system characteristic function under the delay attack is given by

$$P(z) = (h+1)z^{\tau+2} + (-2\eta h\rho + 2\eta h - h - 2\alpha\eta + 2\eta - 1)z^{\tau+1} + 2\alpha\eta h(\rho-1)z^{\tau} + 2\eta h\rho z - 2\alpha\eta h\rho.$$

We numerically compute ROS_{λ_o} for various settings of α and ρ. Fig. 10(b) plots the stability boundaries. We can see that the ROS_{λ_o} expands with α. As we just discussed, a larger α increases the price-elasticity of supply under Case 2. Therefore, it makes the system more capable of responding to generation scheduling errors in time, and hence improves system resilience. Moreover, we can see that when $\rho = 0.5$, the system is stable for $\eta \in (0, 2/3)$. As discussed in Section 6.2.1(3), given the above complex $P(z)$, it is difficult to prove Proposition 5.9 under the stateful supply model. Extensive numerical experiments show that, for $\rho \in (0, 0.5]$, $\tau \in [1, 48]$, and $\alpha \in (0,1)$, $\mathrm{ROS}(\rho,\tau) = \{\eta | 0 < \eta < \frac{1}{1+\alpha}\}$. Since $(0, \frac{1}{1+\alpha})$ is the new stability condition for the pricing algorithm (cf. Proposition 6.4), the numerical results suggest that the memory effect in supply does not change our main result for the delay attack.

6.2.3. Summary. Sections 6.2.1 and 6.2.2 show that the memory effect in supply, which can be observed in real supply-versus-price data traces, does not change the implications of the main results obtained in Section 5, under two extreme approaches to handling the generation scheduling errors. Our analysis can be easily extended to address the case where suppliers and consumers handle complementary portions of the errors. Note that the memory effect could also exist in demand. For instance, the theoretic analysis in [Roozbehani et al. 2012b] based on a load shifting model shows that the aggregated demand is affected by price history. However, modeling stateful effects in demand still remains an open issue. Thus, deviations from the general demand model in Section 3.2 due to various demand-side factors including energy storage and load shifting, as well as the impact of these deviations on our analytic results, are left for future work. However, our analysis in this section that studies the memory effect in supply sheds light on how to address these demand-side factors. Moreover, a set of simulations in Section 8.1.2 show that Proposition 5.9 is still valid when the consumers can shift load according to a simple scheme.

7. DISCUSSIONS

7.1. Impact of Uncertainties in Demand and Supply Models

The previous sections assumed that the supply and price-responsive demand models, i.e., $s(\lambda_k)$ and $w(\lambda_k)$, are deterministic functions of λ_k. In practice, both demand and supply may have uncertainties, such that they will deviate from the deterministic models. This section discusses the impact of these uncertainties on our analytic results. We assume that the demand and scheduled supply are $d(k, \lambda_k) + u_k$ and $s(\lambda_k) + v_k$, respectively, where u_k and v_k are the (price-independent, time-dependent, and bounded) uncertainties in demand and supply, respectively. Thus, we can consider $(v_k - u_k)$ an exogenous bounded disturbance to the system and add it to the disturbance of $(s_o - w_o - b_k)$ shown in Fig. 3. As the analytic results of this paper are based on bounded-input bounded-output stability, as discussed in Section 4, they are still valid in the presence of the exogenous bounded disturbance $(v_k - u_k)$. In Section 8.1, the simulations using uncertain demand that follows a truncated normal distribution validates this discussion.

7.2. Superimposed Attacks, Heterogeneous Attacks, and Other Integrity Attacks

In this section, we discuss how to extend our analysis framework to address a class of integrity attacks that are the superimposition of scaling and delay attacks. We also discuss how to adapt our analysis to scenarios in which the attack models/parameters are different for different compromised consumers. Lastly, we discuss three other integrity attacks that have not been addressed in the previous sections.

From control theory [Ogata 1995], our analysis framework can be applied to derive the ROS_{λ_o} and ROS under any integrity attack that can be modeled as a linear time-invariant (LTI) system with the transfer function $\frac{\Lambda'(z)}{\Lambda(z)} = \frac{\sum_{i=0}^{n} a_i z^{-i}}{\sum_{j=0}^{m} b_j z^{-j}}$, where the $\Lambda(z)$ and $\Lambda'(z)$ are the z-transforms of λ_k and λ'_k. In the time domain, λ'_k is given by the linear combination of λ_{k-i} and λ'_{k-j}, where $0 \leq i \leq n$ and $1 \leq j \leq m$. The scaling and delay attacks are special cases of this general attack model. For instance, under the delay attack, $b_0 = 1$, $b_j = 0$ for $j \geq 1$, $a_\tau = 1$, $a_i = 0$ for $i \neq \tau$. This general attack model can also be regarded as the superimposition of scaling and delay attacks. We now illustrate the enhanced impact of attack superimposition using a simple example: $\lambda'_k = \gamma \cdot \lambda_{k-\tau}$. Under this attack superimposition on the prices for consumers and the stateless supply model, the closed-loop system characteristic function of Eq. (1) is $P(z) = (h+1)z^{\tau+1} + (2\eta + 2\eta(1-\rho)h - h - 1)z^{\tau} + 2\eta\rho h\gamma\mu$, where μ is defined in Section 5.2. We can still apply the Jury test to derive the $\mathrm{ROS}_{\lambda_o}(\rho, \gamma, \tau)$ and $\mathrm{ROS}(\rho, \gamma, \tau)$. Fig. 4(a) shows the stability boundary for this attack superimposition with $\rho = 1$, $\gamma = 0.6$, and $\tau = 1$. The ROS_{λ_o} of this attack superimposition is smaller than the delay attack with $\rho = 1$ and $\tau = 1$, which means stronger attack impact.

If two subsets of consumers are subject to two different attacks that happen simultaneously, Eq. (1) can be rewritten as $w'(\lambda_k) = \rho_1 w(\lambda'_k) + \rho_2 w(\lambda''_k) + (1 - \rho_1 - \rho_2) w(\lambda_k)$, where ρ_1 and ρ_2 are the fractions of consumers subject to the two attacks, and λ'_k and λ''_k are the corresponding compromised prices. Our analysis framework still applies once the models of λ'_k and λ''_k are specified. The attack with different parameters (e.g., consumers are subject to different delays) can be treated as simultaneous attacks.

Lastly, we discuss three integrity attacks that have not been addressed in the previous sections. In a *pulse attack*, *random attack*, and *ramp attack*, a price signal is added with a temporally-spaced bounded pulse signal, a bounded random signal, and a signal that increases or decreases with time over a bounded time period, respectively. The formal definitions of these attacks can be found in [Sridhar and Govindarasu 2014]. Different from the scaling and delay attacks that can be modeled as transfer functions (i.e., blocks in the feedback loop), these three attacks essentially correspond to addition operations with exogenous bounded inputs. From control theory, they will not affect the bounded-input bounded-output stability of the pricing system. But they can cause transient effects that will increase the generation scheduling errors and system operating cost. We now discuss their transient effects. The pulse attack will cause a transient after each pulse, which can be corrected by the feedback-based pricing algorithm. The random attack is an exogenous, bounded, and random disturbance to the system, which increases the system volatility but does not diverge the prices because the pricing system is stable. When the price signals to consumers are under a ramp-up attack, the feedback-based pricing algorithm will continually decrease the price to compensate for the generation scheduling error until the end of the attack.

8. TRACE-DRIVEN SIMULATIONS

We conduct two sets of simulations to evaluate the performance of the RTP algorithm in Eq. (2) and the impact of integrity attacks. The first set of simulations for 1,405 houses conforms to the assumptions in Section 3 and validates the analytic results of this paper. Specifically, we assume that these 1,405 houses are all customers of an RTP system and they are subject to the same real-time price at any time instant. We use GridLAB-D [GridLAB-D 2014], a power distribution network simulator, by extending it to capture the models presented in Section 3. GridLAB-D provides several advantages. First, it captures various realistic factors that are not addressed in our analysis, including physical characteristics of power equipment (e.g., power line capacities and impedances) and power loss. Second, it can record emergency events that occur when the ratings of lines and transformers are exceeded. These events indicate important physical consequences of the integrity attacks such as line trips and service interruptions. The second set of simulations simulates a 4-bus transmission system [Grainger and Stevenson 1994, p. 337] and the IEEE 118-bus system, where the load at each bus is equivalent to that of the 1,405 houses. This set of simulations evaluates the optimality of the RTP algorithm in the absence and presence of attacks, by comparing with the optimal locational marginal prices (LMPs) that are often used to account for locational line losses and congestion in transmission systems. In this section, we focus on integrity attacks on the prices for consumers.

8.1. Simulations for 1,405 Houses

8.1.1. Simulation Methodology and Settings. We use a distribution feeder specification [Schneider et al. 2008] with default settings. This feeder covers a moderately populated urban area and comprises 1,405 houses, 2,134 buses, 3,314 triplex buses, 1,944 transformers, 1,543 overhead lines, 335 underground lines, and 1,631 triplex lines. For this small-scale distribution feeder, LMPs are usually not applicable and hence all the houses are subject to the same price as discussed in Section 3.1. By leveraging the extensibility of GridLAB-D, we develop new modules that implement the CEO model for each single house, the RTP algorithm in Eq. (2), and the attack strategies. We measure the instantaneous power of the entire feeder at the root node. Its peak value over the previous pricing period is used as $d(\lambda_{k-1})$ in Eq. (2). As we focus on evaluating the physical consequences of attacks, we do not simulate the logistics of the attacks and assume that the adversary can gain access to the meters of his choosing. Specifically, if a house is not subject to attacks, it directly reads the real-time price from the *ISO module*; otherwise, it reads the price from an *adversary module* that modifies the price according to the attack models. All the attacks are launched after the system has converged.

We adopt the CEO demand model for each single house, where the parameters are drawn from normal distributions. Specifically, for each consumer i, $D_i \sim \mathcal{N}(6.84, 3.42^2)$ (unit: kW) and $\epsilon_i \sim \mathcal{N}(-0.8, 0.1^2)$. Under this setting, if the price is within $[10, 20]$, the per-house price-responsive demand is within $[0.65, 1.1]$ kW. Moreover, the sum of multiple CEO models is a decreasing function of the price, thus conforming to the assumption in Section 3.2. To improve the realism of the simulations, we use a half-hourly trace of total demand from March 1st to 22nd, 2013, in NSW, Australia (provided by AEMO [Austrilian Energy Market Operator 2014]) as the baseline load. The baseline load of a single house is set to be a scaled version of the real trace data. The resultant range of the per-house baseline load is $[0.276, 0.488]$ kW. Hence, when the price is within $[10, 20]$, the demand of a household is within $[0.9, 1.6]$ kW, which is consistent with the average demand of a household in reality. In our simulations, the price is updated every half an hour, to be consistent with the setting of the demand data traces [Aus-

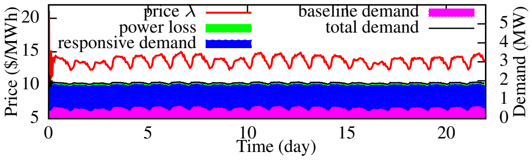

Fig. 11. Price stabilization without attack.

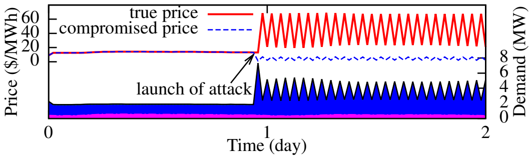

Fig. 12. Scaling attack ($\rho = 65\%$, $\gamma = 0.1$).

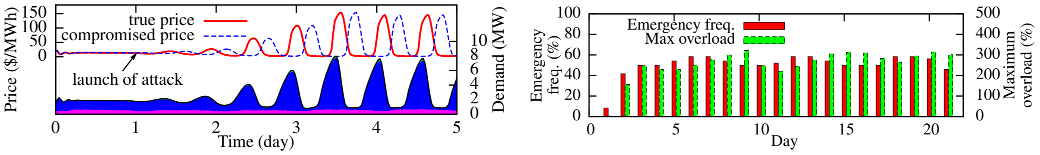

(a) Prices and demands (b) Emergencies

Fig. 13. Impact of delay attack ($\rho = 100\%$, $\tau = 9$).

trilian Energy Market Operator 2014]. In each pricing period, the simulated demand remains constant. For the supply model, the settings obtained in Fig. 1 (i.e., $p = 152$ and $q = 4503$) are for the whole NSW region. They must be scaled down to fit the simulated feeder with 1,405 houses. Since there are 2.8 million households in NSW and 57% of AEMO's supply is for residential demand [Australian Energy Market Operator 2011, p. 15], the two parameters are scaled as follows: $p = \frac{57\% \times 152}{2800000/1405} = 43.638 \times 10^{-3}$ and $q = \frac{57\% \times 4503}{2800000/1405} = 1.287$. Other default settings include: $T = 0.5$, $\lambda_{\min} = 1$, $\lambda_{\max} = 200$, and $\eta = 0.5$.

8.1.2. Simulation Results. The first simulation evaluates the effectiveness of the direct feedback approach [Roozbehani et al. 2012a] and our RTP algorithm in Eq. (2). In the simulations, the direct feedback approach is unstable, where the price oscillates between $\lambda_{\min}$ and $\lambda_{\max}$. The total demand reaches 10 MW a few hours after the start of the simulation, and GridLAB-D reports that four power lines are overloaded. Fig. 11 plots the price and resultant demands under our RTP algorithm in Eq. (2). We can see that the price fluctuates slightly for a few hours after the start of the simulation, due to an inappropriate initial price. After the system converges, it can well adapt to the time-varying baseline load. The generation scheduling error is close to zero, which means that the clearing price is achieved. Moreover, we can see that the power loss is insignificant.

The second simulation evaluates the impact of a scaling attack. Fig. 12 plots the true and compromised prices, as well as the breakdown of demand under the scaling attack. We can see that the price and the demand fluctuates severely. Moreover, the power loss is insignificant. GridLAB-D reports excessive power line overload events after the launch of the attack. We also extensively evaluate the impact of the scaling attack with different settings of ρ and γ. We use the standard deviation of the generation scheduling error after the launch of the attack, denoted by $\sigma(e)$, as the system volatility

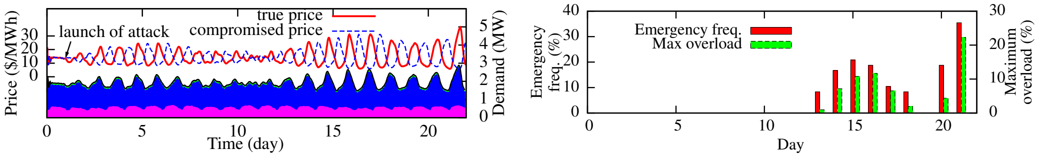

(a) Prices and demands (b) Emergencies

Fig. 14. Impact of delay attack ($\rho = 65\%$, $\tau = 24$).

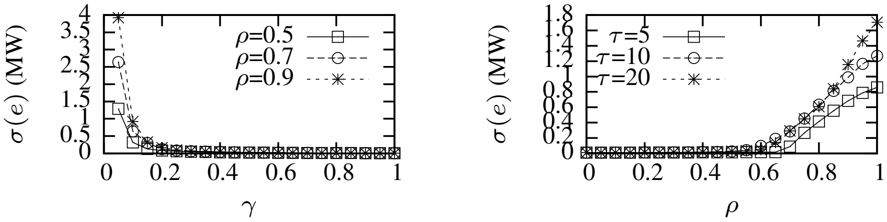

(a) Scaling attack (b) Delay attack

Fig. 15. System volatility under attacks.

metric. A near-zero $\sigma(e)$ means convergence, while a considerably large $\sigma(e)$ means oscillation or divergence. Fig. 15(a) plots $\sigma(e)$ versus γ under various settings of ρ. We can see that the system volatility increases with ρ and decreases with γ.

The third simulation evaluates the impact of delay attacks. Fig. 13(a) and Fig. 14(a) show the evolution of price and the breakdown of demand under the delay attacks with different parameters. We can see that the power loss is insignificant. Thus, it is safe to ignore power loss in the analysis. We also investigate the emergency events reported by GridLAB-D. The *overload* of a line or a transformer is defined as the percentage of the exceeded current/power with respect to the rated value. Fig. 13(b) and Fig. 14(b) plot the *emergency frequency* and maximum overload in each day. The emergency frequency is defined as the ratio of the number of pricing periods with reported emergency events to the number of pricing periods per day (i.e., 48). In Fig. 13, a small generation scheduling error caused by the time-varying baseline load will be amplified iteratively along the control loops, after the launch of the attack. The overload can be up to 350%. In practice, such a high overload will cause circuit breakers to open and hence regional blackouts. In Fig. 14, the system appears to diverge and then converge again without causing any emergencies. However, it diverges again from the 12th day due to the changing baseline load, causing excessive emergencies. This illustrates the stealthiness of the delay attack that causes marginal system stability. We also evaluate the impact of the delay attack with different settings of ρ and τ. The results are shown in Fig. 15(b). We can clearly see that when $\rho < 0.5$, the system remains stable, which is consistent with Proposition 5.9.

The fourth simulation evaluates the impact of demand uncertainties. For each house, in addition to the baseline demand and the price-responsive demand, we generate a random demand every pricing period and add it to the total demand. Denote by u the uncertainty index. The random demand is sampled from a truncated normal distribution with a zero mean and a standard deviation of $\frac{u \cdot b_{k,j}}{3}$, where $b_{k,j}$ is the baseline demand of a household. The random demand is bounded within $[-u \cdot b_{k,j}, u \cdot b_{k,j}]$. Fig. 16 shows the system volatility under different settings for the uncertainty index u and the delay τ. We can see that, without attacks (i.e., $\rho = 0$), the system experiences volatility caused by the demand uncertainty. Moreover, under the two settings for τ in Fig. 16,

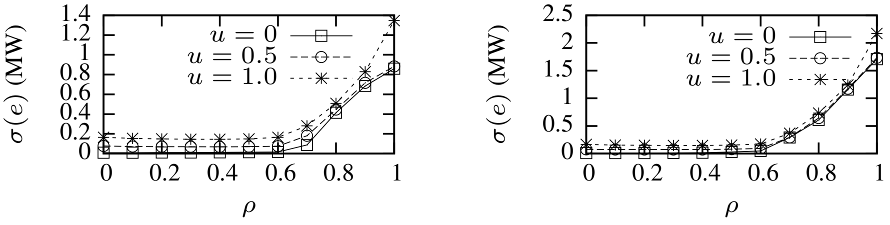

(a) $\tau = 5$

(b) $\tau = 20$

Fig. 16. System volatility under demand uncertainties and delay attacks.

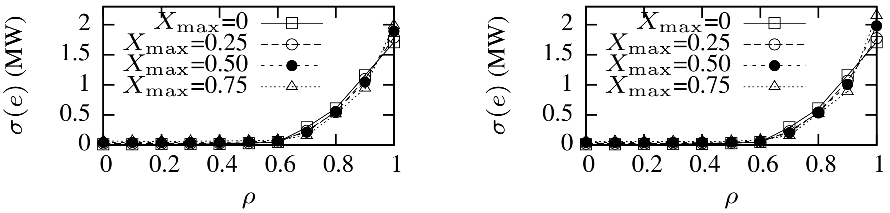

(a) $d_{\max} = 12$

(b) $d_{\max} = 48$

Fig. 17. System volatility with load shifting and under delay attack ($\rho = 1$, $\tau = 20$).

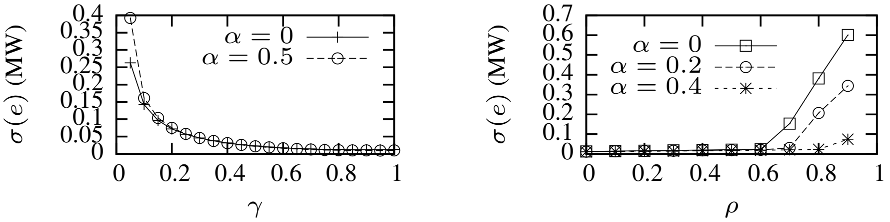

(a) Scaling attack ($\rho = 0.5$)

(b) Delay attack ($\tau = 12$)

Fig. 18. System volatility under stateful suppliers and attacks.

the system volatility exhibits a similar trend as in Fig. 15(b), i.e., $\sigma(e)$ keeps flat when $\rho < 0.5$ and increases drastically with ρ when $\rho > 0.6$.

The fifth simulation evaluates the impact of load shifting. Each consumer follows a load shifting scheme as follows. In the kth pricing period, consumer j shifts $100 \times x\%$ of his price-responsive load, i.e., $x \cdot w_j(\lambda_k)$, to the $(k+d)$th pricing period, where x is a random variable uniformly distributed within $(0, X)$ and d is a random integer uniformly distributed within $[0, d_{\max}]$. To simulate the consumers' incentive to shift load, we update the upper bound X every pricing period by $X = X_{\max} \cdot \frac{\lambda_k - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}}$ such that it is proportional to the price λ_k. Thus, the actual demand of consumer j in the kth pricing period, $d_j(k, \lambda_k)$, is the sum of the baseline demand $b_{k,j}$, the remaining price-responsive load $(1 - x) \cdot w_j(\lambda_k)$, and the sum of previous loads shifted to the kth period. Fig. 17 shows the system volatility under different settings for $X_{\max}$ and $d_{\max}$, when the prices to consumers are under the delay attack with $\rho = 1$ and $\tau = 20$. When $\rho < 0.5$, the load shifting introduces little extra volatility. When ρ is large such that the system is unstable, the system volatility increases with both $X_{\max}$ and $d_{\max}$. This is because, when the price reaches high values in the oscillations, the demand is more uncertain due to larger $X_{\max}$ and $d_{\max}$. Nevertheless, the results in Fig. 17 are still consistent with Proposition 5.9.

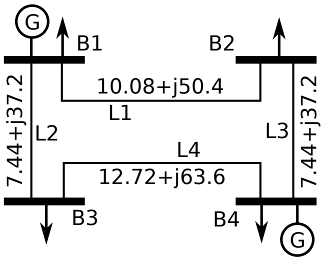

Fig. 19. A 4-bus system [Grainger and Stevenson 1994, p. 337]. Bus voltage is 230 kV.

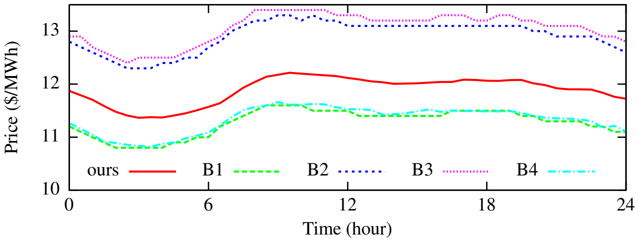

Fig. 20. 4-bus system: the price given by our approach and the optimal LMPs in the 2nd day (base impedance is 529 ohms).

We lastly evaluate the impact of memory effect in the supply. This set of simulations considers Case 2 in Section 6, i.e., the suppliers fully handle the generation scheduling error. From Proposition 6.4, η is set to 0.3 to ensure the stability of the RTP algorithm in the absence of attacks. Fig. 18(a) and 18(b) plot $\sigma(e)$ under the scaling and delay attacks, respectively, with different settings for α (i.e., the weight in Eq. (3)). The $\sigma(e)$ increases or decreases with α under the scaling and delay attacks, respectively, which is consistent with the discussions in Section 6.2.2. Moreover, if $\rho < 0.5$, the system remains stable under the delay attack.

8.2. Simulations for A 4-Bus Transmission System and the IEEE 118-Bus System

8.2.1. Simulation Methodology and Settings. In this section, the first set of simulations is based on the 4-bus transmission system [Grainger and Stevenson 1994, p. 337] shown in Fig. 19. We assume that a load equivalent to 1,405 houses in Section 8.1 is connected to each bus. There are two generators connected to the buses B1 and B4. LMPs are often applied to buses to account for power losses and congestion caused by lines' impedances and limited capacities [Li and Bo 2007]. Our RTP algorithm assumes that the whole system uses the same price, which may lead to a loss of optimality. In this section, we assess the optimality of our algorithm by comparing its results with the corresponding optimal LMPs. We extend an LMP formulation based on the dc optimal power flow (DCOPF) in [Li and Bo 2007] to address a general cost function and price-responsive demand. The details of the formulation are omitted due to space constraints and can be found in Appendix B of the supplementary file of this paper. Note that the optimal solution to LMP problem cannot be known if the ISO has limited knowledge about the demand models and value functions. Under our RTP algorithm, the generators at bus B4 and bus 69 in the 4-bus and 118-bus systems, respectively, are designated to compensate for the generation scheduling error such that the power balance is satisfied. We set large capacities for all the branches such that they are not congested. This setting allows us to evaluate the required line capacity under attacks. The second set of simulations is based on the IEEE 118-bus system.

Table II. 4-bus system: optimality of our control-theoretic RTP algorithm.

Base impedance	Our RTP		Optimal LMPs					$\Delta\omega$
(ohm)	λ	Loss	λ_1^B	λ_2^B	λ_3^B	λ_4^B	Loss	($/hour)
529^a	12.2	3.78%	11.6	13.3	13.4	11.6	3.45%	0.18
2×529	12.7	7.19%	11.4	14.8	15.0	11.8	6.11%	0.63
4×529	13.7	13.5%	10.8	17.6	17.9	12.3	10.1%	2.36
6×529	15.0	20.2%	10.4	20.4	20.7	12.8	13.0%	6.03
8×529	17.4	30.2%	9.93	23.1	23.4	13.3	15.2%	16.8

Note: This table reports average results over 22 days.
[a]The value 529 ohms is from [Grainger and Stevenson 1994, p. 337].

Table III. 4-bus system: the power flows under attacks.

Attack	Algorithm	F_1	F_2	F_3	F_4	Loss	$\Delta\omega$
No	Our RTP	0.79	1.15	1.03	0.66	3.78%	0.18
	Optimal LMPs	0.74	1.08	0.96	0.61	3.45%	n/a
Scaling attack (γ=0.1)	Our RTP	3.25	3.43	4.08	3.01	14.1%	316
	LMPs with partial knowledge	4.04	1.43	12.19	9.49	40%	13304
	LMPs with full knowledge	0.64	0.89	0.83	0.54	1.42%	n/a
Delay attack (τ=4)	Our RTP	4.80	5.34	7.98	6.12	15.7%	n/a
	LMPs with partial knowledge	0.77	1.09	1.09	0.72	3.45%	n/a

Note: Attacks are launched at the start of the 2nd day. The power flows (unit: MW) in this table are the maximums of abstract values over 20 days from the 3rd day. The base impedance is 529 ohms and $\rho = 1$. The $\Delta\omega$ is with respect to the ω achieved by the LMPs with full information.

8.2.2. Simulation Results for the 4-Bus System. First, we evaluate the optimality of our RTP algorithm in the absence of attacks. Fig. 20 shows the price given by our algorithm and the optimal LMPs on the second day. As the buses B2 and B3 draw power from B1 and B4, they need to pay the costs for the power losses in transmission. Thus, their LMPs are higher than those of B1 and B4. Table II shows statistical results for the prices and total power losses, under different settings of base impedance. Note that the impedance of a branch is the product of the per-unit value in Fig. 20 and the base impedance. The power losses in Table II are the ratios of total power loss to the total supply. When the power loss is below 10%, which is commonly seen in practice (e.g., 7% in U.S. [U.S Energy Information Administration 2014]), the power loss under our approach is at most 1% more than that under the optimal LMPs. Thus, in the absence of attack, our algorithm yields near-optimal performance. Table II also lists the reduction of social welfare, denoted by $\Delta\omega$, with respect to that of LMPs.

Second, we evaluate the impact of attacks on the 4-bus transmission system. Table III lists the power flows on the four branches in the absence versus presence of attacks. When solving the LMPs in the presence of attacks, we make different assumptions regarding the ISO's knowledge. By *partial knowledge*, we mean that the ISO knows the value functions and demand models of the consumers. By *full knowledge*, we mean that the ISO further knows the attack parameters (e.g., ρ and γ). Note that, as discussed in Section 3.2, in practice, it is difficult to know the consumer models and attack parameters. The LMP results only help understand the performance limit. Under the scaling attack, the prices under our RTP approach oscillate within the range $[30, 80]$ and the power flows increase, causing a power loss rate of up to 14%. Under the LMP approach with partial knowledge, the scaling attack will significantly increase the power flows and loss rate, because we lack a feedback mechanism to correct the large generation scheduling errors and the associated high power flows caused by the increased demand. As it is extremely difficult to solve the LMPs with full knowledge under the delay attack since the social welfares in different pricing periods are correlated, we skip the evaluation of LMPs with full knowledge under the delay attack. From Table III, under our RTP approach, the delay attack results in increased power

Table IV. 118-bus system: Average power loss rates without or under attacks.

Attack	Algorithm	Loss
No	Our RTP	4.8%
	LMPs	4.0%
Scaling attack (γ=0.4)	Our RTP	6.2%
	LMPs with full knowledge	5.2%
Delay attack ($\tau = 4$)	Our RTP	9.5%

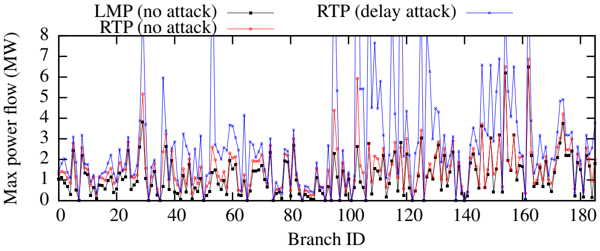

Fig. 21. 118-bus system: all branches' maximum power flows.

flows and loss rates due to price oscillations caused by instability. Under the LMP approach with partial knowledge, which has no closed loops and thus no stability issues, the delay attack introduces little impact to power flows and loss rates. However, as just discussed, this LMP approach cannot be implemented in practice. In summary, under our RTP approach, the price oscillations caused by attacks can cause significantly increased power flows and power losses on transmission lines in a physical system.

8.2.3. Simulation Results for the IEEE 118-Bus System. We also evaluate our RTP algorithm using the IEEE 118-bus system. Table IV summarizes the average power loss rates under the different approaches. Similar to the results for the 4-bus system, our algorithm yields near-optimal performance in the absence of attacks. In the presence of a scaling attack with $\gamma = 0.4$, the RTP system under our approach is still stable. Thus, our approach yields a power loss rate close to the LMP approach with full knowledge. We note that, if $\gamma = 0.1$ as in Section 8.2.2, the price will diverge, leading to a high power loss rate. In the presence of a delay attack, the power loss rate increases to 9.5%. Fig. 21 plots the maximum power flow of each branch in the simulation. We can see that in the presence of the delay attack, the branch power flows increase significantly, resulting in the high power loss rate in Table IV.

8.2.4. Discussion. As stated in Section 3.1, this paper assumes that the transmission lines are not congested. We now use the scaling attack on consumers' price signals as an example to discuss its impact on a power grid when some transmission lines are already congested before the launch of the attack. Before the attack, the solution to the LMP problem fully utilizes the capacities of some of the transmission lines (i.e., these lines are congested). When the consumers' price signals are under a scaling attack with $\gamma < 1$, the consumers will increase demand, causing increased line flows as observed in Table III. Depending on which consumers' price signals are compromised, the attack may increase the power flows of the congested lines beyond the line capacities. In practice, the overloads may lead to line trips and cascading failures. Note that our analysis of the stability of closed-loop RTP systems is applicable before any line trips and cascading failures happen.

Table V. Summary of Assumptions

Section		Demand model	Supply model	Pricing model
§4	§4.1	general demand model in §3.2	general supply model in §3.3	direct feedback
	§4.2	linearization to the general demand model in §3.2	linearization to the general supply model in §3.2	Proposition 4.1
§5	§5.1-§5.3			
	§5.4	CEO demand model	linear & quadratic supply models	Eq. (2)
§6	§6.1-§6.2	linearization to the general demand model in §3.2	Eq. (3), where $s(\lambda_k)$ is given by the linearized supply model	Proposition 4.1
§8	§8.1	CEO demand model with demand uncertainties & load shifting	linear supply model & Eq. (3)	Eq. (2) & direct feedback
	§8.2	CEO demand model	linear supply model	Eq. (2) & LMP

9. CONCLUSION AND DISCUSSIONS

This paper investigates the impact of scaling and delay attacks on the stability of RTP systems. The analysis and results of this paper are based on a basic set of assumptions presented in Section 3 and necessary variations made in Section 4 to Section 8. These variations are summarized in Table V. We characterize the impact using a control-theoretic metric, namely the *region of stability*. We show that, to destabilize the RTP system, it is necessary for the adversary to reduce the prices for consumers in a scaling attack or compromise more than half of the prices for consumers in a delay attack. We conduct trace-driven simulations to validate our analysis. The results of this paper improve our understanding of the security of RTP systems so that suitable defensive measures can be taken in response.

9.1. Limitations, Discussions, and Future Work

In this paper, we have made several simplifying assumptions that enable us to focus on the essence of the problem. This section summarizes these simplifications and discusses future work directions to address them. These simplifications can be summarized into the following aspects:

Simplified demand and supply models. The deterministic demand and supply models in Sections 3.2 and 3.3 do not capture several realistic factors such as uncertainties in consumer and supplier behaviors, generation ramp constraints, load shifting, and energy storage. In this paper, we have discussed and evaluated the potential impacts of some of these factors on our results. Specifically, we have discussed the impact of bounded additive deviations from the models (Section 7.1) and evaluated the impact of specific types of random demand and load shifting through simulations (Section 8.1). We have also incorporated stateful effect in supply, which can be a result of generation ramp constraint (Section 6). However, further analysis is still required to explicitly address these affecting factors, by integrating new results on a few open and active research topics (e.g., demand-side energy storage and load shifting).

Linearization and approximation. The results of this paper are obtained under linearized demand and supply models as well as an approximation of total demand when a fraction of consumers receive compromised prices (Eq. (1)). The linearization will lead to inaccuracies in characterizing systems with non-linear supply and/or demand models. We expect our linearization-based analysis provides insights into understanding the problem under non-linear settings, as well as a basis for finer analysis based on piece-wise linearization that can be a focus of future work.

Simplified market model and simple RTP algorithm. This paper adopts a simplified market model that preserves the principle of RTP, i.e., the real-time wholesale price is directly relayed to end consumers. This model does not address various practical factors such as bidding markets, advance notice of price, and ex-post price adjustment. Future research is still needed to incorporate these factors into the analysis. The design of RTP algorithms is still under active research and an academic consensus has not been achieved. The security analysis and results in this paper are based on a simple RTP algorithm (Proposition 4.1), which is similar to the one in [Dalkilic et al. 2013]. We note that our and Dalkilic's algorithms are different from those adopted in the current wholesale markets, which often solve the clearing prices based on autoregression-based demand predictors. As discussed in Section 4.1 and [Roozbehani et al. 2012a], the stability of these autoregression-based RTP algorithms still needs further careful examinations and may need substantial retrofits when more end consumers are exposed to real-time prices. Thus, the results of this paper cannot be directly applied to the current wholesale markets using autoregression-based RTP algorithms, since we focus on the scenario where all end consumers are exposed to real-time prices.

No transmission system constraints and identical prices across locations. Our analysis does not account for specific affecting factors of transmission systems that are often addressed by locational prices. Such factors include transmission cost and line congestion. Because of this simplification, we assume that the real-time prices are identical across all locations. Although in Section 8.2.4 we use an example to discuss the attack impact on a congested transmission system, the baseline design of stable locational RTP algorithms and whether our results still hold for these algorithms remain open and interesting research problems. The simulations in Section 8.1 based on the distribution system simulator, GridLAB-D, provide a validation for our analysis that ignores transmission system constraints. To validate the analytic results that address transmission system constraints, GridLAB-D will not be appropriate.

Specific attack models. The results of this paper are specific to two attack models (scaling and delay attacks). In Section 7.2, we have also discussed three other integrity attacks (pulse, random, and ramp attacks). All these attacks discussed in this paper follow certain rules to change the price signals. They may be accomplished by indirect techniques that are less effort-intensive. It is also interesting to study the impact of the attacks that do not have to conform to certain rules, if the attacker gains the privilege of directly changing the price signals.

Others. The following aspects are worth further studies. First, our analysis assumes that the ISO can estimate $\dot{\omega}(\lambda_o)$ accurately. A bounded estimation error for $\dot{\omega}(\lambda_o)$ can be integrated into the analysis. Second, in addition to stability that is the focus of this paper, the transient of an RTP system under attack is also an important aspect. Third, it is interesting to develop attack detection algorithms based on those proposed for general closed-loop systems (e.g., [Eyisi and Koutsoukos 2014]).

REFERENCES

Hunt Allcott. 2009. Real Time Pricing and Electricity Markets. (2009). http://economics.stanford.edu/files/Allcott3_13.pdf.

Fernando Alvarado. 1999. The Stability of Power System Markets. *IEEE Transactions on Power Systems* 14, 2 (1999), 505–511.

Ameren. 2014. Real-Time Pricing for Residential Customers. (2014). http://www.ameren.com/sites/aiu/ElectricChoice/Pages/ResRealTimePricing.aspx.

Saurabh Amin, Xavier Litrico, Shankar Sastry, and Alexandre M. Bayen. 2013. Cyber Security of Water SCADA Systems-Part I: Analysis and Experimentation of Stealthy Deception Attacks. *IEEE Transactions on Control Systems Technology* 21, 5 (2013), 1963–1970.

Australian Energy Market Operator. 2011. 2011 National electricity forecasting. (2011). http://www.aemo.com.au/~/media/Files/Other/forecasting/0400-0053%20pdf.pdf.

Austrilian Energy Market Operator. 2014. (2014). http://www.aemo.com.au.

Galen Barbose, Charles Goldman, Ranjit Bharvirkar, Nicole Hopper, Michael Ting, and Bernie Neenan. 2005. *Real Time Pricing as a Default or Optional Service for C&I Customers: A Comparative Analysis of Eight Case Studies*. Technical Report. Lawrence Berkeley National Laborary.

Alvaro A. Cárdenas, Saurabh Amin, Zong-Syun Lin, Yu-Lun Huang, Chi-Yen Huang, and Shankar Sastry. 2011. Attacks against Process Control Systems: Risk Assessment, Detection, and Response. In *6th ACM*

Symposium on Information, Computer and Communications Security (ASIACCS). ACM, Hong Kong, 355–366.

Alvaro A. Cárdenas, Saurabh Amin, and Shankar Sastry. 2008. Secure Control: Towards Survivable Cyber-Physical Systems. In *28th International Conference on Distributed Computing Systems Workshops*. IEEE, Beijing, China, 495–500.

Joon Young Choi, Seong-Hwang Rim, and Jong-Keun Park. 1998. Optimal Real Time Pricing of Real and Reactive Powers. *IEEE Transactions on Power Systems* 13, 4 (1998), 1226–1231.

ComEd. 2014. ComEd Residential Real-Time Pricing Program. (2014). https://rrtp.comed.com/.

Ozgur Dalkilic, Atilla Eryilmaz, and Xiaojun Lin. 2013. Stable Real-Time Pricing and Scheduling for Serving Opportunistic Users with Deferrable Loads. In *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, Monticello, Illinois, USA, 1200–1207.

Mike Davis. 2009. Recoverable Advanced Metering Infrastructure. (2009). Black Hat Technical Security Conference, Las Vegas, Nevada, USA.

Emeka Eyisi and Xenofon Koutsoukos. 2014. Energy-Based Attack Detection in Networked Control Systems. In *The 3rd international conference on High confidence networked systems (HiCoNS)*. ACM, ACM, Berlin, Germany, 115–124.

Massimo Filippini. 2011. Short- and Long-Run Time-of-Use Price Elasticities in Swiss Residential Electricity Demand. *Energy Policy* 39, 10 (2011), 5811–5817.

Stein-Erik Fleten and Erling Pettersen. 2005. Constructing Bidding Curves for a Price-Taking Retailer in the Norwegian Electricity Market. *IEEE Transactions on Power Systems* 20, 2 (2005), 701–708.

Georgia Power. 2014. RTP-HA-2 Program. (2014). http://www.georgiapower.com.

John J. Grainger and William D. Stevenson. 1994. *Power System Analysis*. McGraw-Hill, New York.

GridLAB-D. 2014. GridLAB-D. (2014). http://www.gridlabd.org.

Liyan Jia, Robert J. Thomas, and Lang Tong. 2012. Impacts of Malicious Data on Real-Time Price of Electricity Market Operations. In *45th Hawaii International Conference on System Sciences (HICSS)*. IEEE, Maui, Hawaii, USA, 1907–1914.

Oliver Kosut, Liyan Jia, Robert J. Thomas, and Lang Tong. 2011. Malicious Data Attacks on the Smart Grid. *IEEE Transactions on Smart Grid* 2, 4 (2011), 645–658.

Fangxing Li and Rui Bo. 2007. DCOPF-based LMP Simulation: Algorithm, Comparison with ACOPF, and Sensitivity. *IEEE Transactions on Power Systems* 22, 4 (2007), 1475–1485.

Na Li, Lijun Chen, and Steven H Low. 2011. Optimal Demand Response Based on Utility Maximization in Power Networks. In *IEEE Power and Energy Society General Meeting*. IEEE, Detroit, Michigan, USA, 1–8.

Mark G. Lijesen. 2007. The Real-Time Price Elasticity of Electricity. *Energy Economics* 29, 2 (2007), 249–258.

Jie Lin, Wei Yu, Xinyu Yang, Guobin Xu, and Wei Zhao. 2012. On False Data Injection Attacks against Distributed Energy Routing in Smart Grid. In *3rd International Conference on Cyber-Physical Systems (ICCPS)*. ACM, Beijing, China, 183–192.

Yao Liu, Peng Ning, and Michael K. Reiter. 2011. False Data Injection Attacks Against State Estimation in Electric Power Grids. *ACM Transactions on Information and System Security* 14, 1 (2011), 13:1–13:33.

Stephen McLaughlin, Dmitry Podkuiko, Sergei Miadzvezhanka, Adam Delozier, and Patrick McDaniel. 2010. Multi-Vendor Penetration Testing in the Advanced Metering Infrastructure. In *Annual Computer Security Applications Conference (ACSAC 26)*. ACM, Orlando, Florida, USA, 107–116.

Tyler Nighswander, Brent Ledvina, Jonathan Diamond, Robert Brumley, and David Brumley. 2012. GPS Software Attacks. In *19th ACM Conference on Computer and Communications Security (CCS)*. ACM, Raleigh, North Carolina, USA, 450–461.

Katsuhiko Ogata. 1995. *Discrete-Time Control Systems* (2 ed.). Prentice-Hall, Englewood Cliffs, New Jersey.

Oracle. 2012. Multiple Vulnerabilities in NTP. (2012). https://blogs.oracle.com/sunsecurity/entry/multiple_vulnerabilities_in_network_time.

PJM. 2014. PJM Operational Data. (2014). http://www.pjm.com/pub/account/lmpgen/lmppost.html.

Public Act 094-0977. 2014. (2014). http://www.ilga.gov/legislation/publicacts/94/PDF/094-0977.pdf.

Li Ping Qian, Ying Jun Angela Zhang, Jianwei Huang, and Yuan Wu. 2013. Demand Response Management via Real-Time Electricity Price Control in Smart Grids. *IEEE Journal on Selected Areas in Communications* 31, 7 (2013), 1268–1280.

Mardavij Roozbehani, Munther A. Dahleh, and Sanjoy K. Mitter. 2012a. Volatility of Power Grids under Real-Time Pricing. *IEEE Transactions on Power Systems* 27, 4 (2012), 1926–1940.

Mardavij Roozbehani, M Ohannessian, D Materassi, and M. A. Dahleh. 2012b. Load-Shifting under Perfect and Partial Information: Models, Robust Policies, and Economic Value. (2012). https://dahleh.lids.mit.edu/wp-content/uploads/2012/08/2012-Load-ShiftingUnderPerfectAndPartialInformation.pdf.

Ishtiaq Rouf, Hossen Mustafa, Miao Xu, Wenyuan Xu, Rob Miller, and Marco Gruteser. 2012. Neighborhood Watch: Security and Privacy Analysis of Automatic Meter Reading Systems. In *19th ACM Conference on Computer and Communications Security (CCS)*. ACM, Raleigh, North Carolina, USA, 462–473.

Pedram Samadi, A-H Mohsenian-Rad, Robert Schober, Vincent WS Wong, and Juri Jatskevich. 2010. Optimal Real-Time Pricing Algorithm Based on Utility Maximization for Smart Grid. In *1st IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, Gaithersburg, Maryland, USA, 415–420.

Pedram Samadi, Hamed Mohsenian-Rad, Vincent WS Wong, and Robert Schober. 2014. Real-Time Pricing for Demand Response Based on Stochastic Approximation. *IEEE Transactions on Smart Grid* 5, 2 (2014), 789–798.

Kevin P. Schneider, Yousu Chen, David P. Chassin, Robert Pratt, Dave Engel, and Sandra Thompson. 2008. Modern Grid Initiative Distribution Taxonomy Final Report. (2008).

Schneider Electric. 2014. Technical Note of ION Smart Meter. (2014). http://www.powerlogic.com/literature/ION%20Time%20Synchronization%20and%20Timekeeping.pdf.

Shams N Siddiqi and Martin L Baughman. 1993. Reliability Differentiated Real-Time Pricing of Electricity. *IEEE Transactions on Power Systems* 8, 2 (1993), 548–554.

Siddharth Sridhar and Manimaran Govindarasu. 2014. Model-Based Attack Detection and Mitigation for Automatic Generation Control. *IEEE Transactions on Smart Grid* 5, 2 (2014), 580–591.

James L. Sweeny. 2002. *The California Electricity Crisis*. Hoover Institution Press, Stanford, California, USA.

Symantec. 2014. Symantec Security Response. Dragonfly: Cyberespionage Attacks Against Energy Suppliers. (2014). http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/Dragonfly_Threat_Against_Western_Energy_Suppliers.pdf.

The Wall Street Journal. 2009. Spies Penetrate U.S. Electric Grid. (2009). http://www.wsj.com/articles/SB123914805204099085.

U.S Energy Information Administration. 2014. How Much Electricity is Lost in Transmission and Distribution in U.S.? (2014). http://www.eia.gov/tools/faqs/faq.cfm?id=105&t=3.

Le Xie, Yilin Mo, and Bruno Sinopoli. 2011. Integrity Data Attacks in Power Market Operations. *IEEE Transactions on Smart Grid* 2, 4 (2011), 659–666.

Rongshan Yu, Wenxian Yang, and Susanto Rahardja. 2012. A Statistical Demand-Price Model with Its Application in Optimal Real-Time Price. *IEEE Transactions on Smart Grid* 3, 4 (2012), 1734–1742.

Yanling Yuan, Zuyi Li, and Kui Ren. 2011. Modeling Load Redistribution Attacks in Power Systems. *IEEE Transactions on Smart Grid* 2, 2 (2011), 382–390.