

Impact of Integrity Attacks on Real-Time Pricing in Smart Grids

Rui Tan¹ Varun Badrinath Krishna¹ David K. Y. Yau^{1,2} Zbigniew Kalbarczyk³

¹Advanced Digital Sciences Center, Illinois at Singapore

²Singapore University of Technology and Design, Singapore

³University of Illinois at Urbana-Champaign, Urbana, IL, USA

{tanrui, varun.bk, david.yau}@adsc.com.sg, kalbarcz@illinois.edu

ABSTRACT

Modern information and communication technologies used by smart grids are subject to cybersecurity threats. This paper studies the impact of integrity attacks on real-time pricing (RTP), a key feature of smart grids that uses such technologies to 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 inefficiency and even severe failures such as blackouts. This paper adopts a control-theoretic approach to deriving the fundamental conditions of RTP stability under two broad classes of integrity attacks, namely, the *scaling* and *delay* attacks. We show that the RTP system is at risk of being destabilized only if the adversary can compromise the price signals advertised to smart meters by 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

Keywords

Smart grid; real-time pricing; stability; cyber security

1. INTRODUCTION

A smart grid is an enhanced electrical grid that uses modern information and communication technologies to improve

system reliability and efficiency. However, these computerized and networking technologies are subject to security threats that range from personal breaches [22] to sophisticated cyber attacks launched by hostile organizations to cause widespread outages [31]. As a sophisticated cyber-physical system, a smart grid features complex closed-loop feedback controls in various physical [18] and economic components [7], which maintain desirable system performance in the presence of dynamics and uncertainties. However, the impacts of cyber attacks against these closed loops on smart 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) in smart grids, which involves closed-loop controls to stabilize the electricity market.

Dynamic pricing [10] is a widely adopted means to balance electricity generation and consumption. The electricity price in the wholesale market is updated periodically (e.g., every hour) 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 users, RTP has been considered a key feature of smart grids, which can reduce over-provisioning and improve system efficiency. Unfortunately, as analyzed in [25], 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 [25]. 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CCS'13, November 4–8, 2013, Berlin, Germany.

Copyright 2013 ACM 978-1-4503-2477-9/13/11 ...\$15.00.

<http://dx.doi.org/10.1145/2508859.2516705>.

be iteratively amplified by the feedback, causing inefficiency and significantly fluctuating demand, and possibly leading to severe failures such as blackouts.

In smart 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. In particular, recent studies [22, 26] have shown that many smart meters lack basic security measures to ensure the integrity and authenticity of their input/output data. 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 severe failures such as outages? 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 attack 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 real-time price control 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 baseline understanding of the security of these systems. We adopt a control-theoretic metric, namely, the *region of stability*, to characterize the security 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, where the prices advertised to smart meters are compromised by a scaling factor (so that the meters will use the wrong prices) and by corrupted timing information (so that the meters will use old prices), respectively. In addition to directly tampering with traffic sent to the smart meters, 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 deployed smart meters. Note that current commercial smart meters [28] synchronize their clocks by either built-in GPS receivers or a network time protocol (NTP) supported by time servers [28]. Both approaches have been shown to be vulnerable to realistic attack methods [4, 23]. As both attacks can be modeled as LTI transfer functions, the security analysis can be conducted under a LTI setting.

Based on our analytical framework, we derive the region of stability for both the scaling and delay attacks. In particular, we show that the RTP system will remain stable if (i) the compromised prices are amplified versions of the true prices under the scaling attack (i.e., the scaling fac-

tor exceeds one), 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 severely fluctuating demand arising from the system instability. We report GridLAB-D [2] simulation results for a distribution system consisting of 1405 consumers to verify our analysis and demonstrate possible system emergencies (e.g., line and transformer overload events) caused by the integrity attacks. Our results provide insights for securing RTP in smart grids. For instance, for the adversary to achieve the goal of compromising the price signals to at least half of the consumers, she may focus her efforts on shared support infrastructures such as the NTP time servers. This highlights the importance of securing these 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. Section 5 analyzes the impact of the scaling and delay attacks on the RTP system. Section 6 discusses extensions of the analysis framework to address a broader class of attacks that are combinations of multiple scaling and delay attacks. Section 7 presents simulation results. Section 8 concludes.

2. RELATED WORK

The security of smart grids is attracting increasing research attention. In particular, false data injection attacks against the state estimation of electrical grids have been extensively studied. In [21], Liu et al. systematically examine the conditions for bypassing a bad data detection mechanism of state estimation under various adversary capability models. Later studies [16, 17, 20, 32, 33] show that the false data injection attacks can lead to increased system operation costs due to inordinate generation dispatch [33] or energy routing [20], as well as economic losses due to misconduct of electricity markets [16, 17, 32]. In particular, the studies in [16, 17, 32] focus on false data injection attacks on real-time wholesale markets. They primarily emphasize attacks on critical measurements, which are often well protected by system operators. Moreover, they ignore demand response of end users to prices. In contrast, our work considers integrity attacks that target distributed smart meters that are relatively much more vulnerable, and also accounts for demand response involving the end users. All these related studies [16, 17, 20, 21, 32, 33] analyze attacks on systems using constrained optimization formulations such as power flow dispatch. The closed loop characterizing the RTP system in our work imposes specific challenges in the security analysis due to its iterative nature.

The security of a broader class of cyber-physical systems that feature complex closed loops has been studied recently. In [12], Cárdenas et al. identify challenges in the security analysis of these systems. In [11], 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 [8], 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. They fall short of characterizing fully the fundamental critical stability conditions.

Table 1: Summary of Notation*

Symbol	Definition	Unit
T	pricing period	hour
k	index of current pricing period	n/a
λ_k	true price signal, $\lambda_k \in [\lambda_{\min}, \lambda_{\max}]$	\$/MWh
λ_k^c	compromised price signal	\$/MWh
λ_k^*	clearing price	\$/MWh
b_k	total baseline demand	MW
$w(\lambda_k)$	total price-responsive demand	MW
$d(k, \lambda_k)$	total demand	MW
D	a constant in the CEO model	MW
ϵ	price elasticity of demand	n/a
$s(\lambda_k)$	scheduled total generation	MW
$\hat{s}(\lambda_k)$	realized total generation	MW
p	slope of linear supply model	MW/\$/MWh
q	intercept of linear supply model	MW
e_k	generation scheduling error	MW
η	price stabilization gain, $\eta \in (0, 1)$	n/a
λ_o	operating point of price stabilization	\$/MWh
C	the set of all consumers	n/a
C'	the set of consumers under attack	n/a
ρ	$\simeq C' / C $	n/a
γ	amplification of scaling attack	n/a
τ	time delay of delay attack	T
h	marginal demand-supply ratio	n/a

* The unit of a quantity is omitted in the paper if it is specified here.

3. PRELIMINARIES

This section presents the market model adopted in this paper, which comprises an *independent system operator* (ISO) (Section 3.1), a set of consumers (Section 3.2), and a set of suppliers (Section 3.3). For both consumers and suppliers, we first describe general abstract models, and then discuss concrete empirical models commonly used in literature. The analytical results in this paper (i.e., Propositions 1 to 5) are based on the abstract models, while the empirical models (i.e., a constant elasticity of own-price demand model and a linear supply model) are used for the numerical examples and simulations. The notation used in this paper is summarized in Table 1. We also use the following mathematical notation: $\dot{f}(x)$ denotes the first derivative of function $f(x)$; $f^{-1}(x)$ denotes the inverse of function $f(x)$; $\mathbb{R}^+/\mathbb{R}^-$ denotes the set of positive/negative real numbers; \mathbb{Z}^+ denotes the set of positive integers.

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 end consumers. Specifically, the price for the k th pricing period $[k \cdot T, (k+1) \cdot T]$, denoted by λ_k , is announced at time instant $(k - k_0) \cdot T$, where k_0 is a non-negative integer. 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 (e.g., [25] and references therein) and is consistent with the essence of several experimental RTP programs [10] provided by utilities, which include Board of Public Utilities in New Jersey, Baltimore Gas and Electric Company in Maryland, and Duquesne Light in Pennsylvania. In these

programs, the hourly wholesale prices published by PJM Interconnection LLC¹ are used directly as retail prices, where $T = 1$ and $k_0 = 0$. A few other experimental RTP programs give customers advance notice of hourly prices. For instance, $k_0 = 1$ for the RTP-HA-2 program of Georgia Power [6]. To simplify the discussion, we focus on RTP schemes without advance notice, i.e., $k_0 = 0$. However, our analysis can be easily extended to encompass advance notice. In reality, locational prices can be applied to address location-dependent transmission costs. In many areas, as generation cost dominates transmission cost, variations of locational prices are often small. For instance, the relative standard deviation of the locational prices for 219 locations published by PJM is often around 5% only [5]. 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 λ_k .

3.2 Consumers

Abstract demand model: Let C denote the set of consumers in the system. For consumer $j \in C$, let $b_{k,j}$ denote the *baseline demand* in the k th pricing period, which is 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. Let $v_j(x - b_{k,j})$ denote the additional value (unit: \$/hour) derived from consuming a total of x units of power in the k th pricing period, where $x \geq b_{k,j}$. We assume that $v_j(\cdot)$ is a strictly increasing and strictly concave function. Let $d_j(k, \lambda_k)$ denote the demand of consumer j in the k th pricing period given price λ_k . We denote $w_j(\lambda_k) = \dot{v}_j^{-1}(\lambda_k)$, where $\dot{v}_j^{-1}(\cdot)$ represents the inverse function of $\dot{v}_j(\cdot)$. The demand is given by

$$d_j(k, \lambda_k) = \arg \max_x (v_j(x - b_{k,j}) - \lambda_k \cdot x) = b_{k,j} + w_j(\lambda_k).$$

In the above equation, $v_j(x - b_{k,j}) - \lambda_k \cdot x$ is the additional utility beyond the baseline by consuming totally x units of power. The $w_j(\lambda_k)$ is referred to as *price-responsive demand*. It is easy to verify that $w_j(\lambda_k)$ is a decreasing function of λ_k . 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, denoted by $d(k, \lambda_k)$, is given by

$$d(k, \lambda_k) = \sum_{j \in C} d_j(k, \lambda_k) = b_k + w(\lambda_k). \quad (1)$$

As there are a large number of consumers, we assume that $b_{k,j}$ and $v_j(\cdot)$ are unknown to the ISO. However, the ISO knows the historical total demand $\{d(h, \lambda_h) | h \in [0, k-1]\}$. The above derivations, which are based on the basic concept of *utility* in economics, explain the consumer's demand response to price. Human-induced demand response has been observed in previous studies [29]. With the increasing adoption of smart appliances and home automation systems, this demand response will become more automated.

Empirical demand model: The constant elasticity of own-price (CEO) model [14] is a simple model that can be used to characterize the total price-responsive demand, which is defined by $w(\lambda_k) = D \cdot \lambda_k^\epsilon$, where D and ϵ are positive and negative constants, respectively. The ϵ is referred to as the *price elasticity of demand*, which is typically within $(-1, 0)$ [13, 19].

¹PJM is a Regional Transmission Organization (RTO). This paper does not distinguish between ISO and RTO.




Figure 1: Total supply vs. wholesale price [1].

3.3 Suppliers

Abstract supply model: Each supplier aims to maximize its profit. Let S denote the set of suppliers in the system. For any supplier $i \in S$, let the function $c_i(x)$ represent the cost (unit: \$/hour [15, p. 534]) of producing and transmitting x units of power. We assume that $c_i(x)$ is a strictly convex and non-negative function over the support $x \geq 0$. Moreover, we assume that $c_i(x)$ is an asymptotically increasing function, i.e., $\exists x_0 \geq 0, \dot{c}_i(x) > 0$ if $x > x_0$. Let $s_i(\lambda_k)$ denote the quantity of power that supplier i schedules to generate in the k th pricing period given price λ_k , which is given by $s_i(\lambda_k) = \arg \max_x (\lambda_k \cdot x - c_i(x)) = \dot{c}_i^{-1}(\lambda_k)$. Note that $\lambda_k \cdot x - c_i(x)$ is the profit from generating x units of power. It is easy to verify that $s_i(\lambda_k)$ is an increasing function of λ_k . We assume that the generation capacity of the supplier i is at least $s_i(\lambda_{\max})$. In this paper, we consider centralized bulk generation rather than distributed generation. Therefore, an ISO can estimate $\{c_i(\cdot) | \forall i \in S\}$ as there are typically a limited number of suppliers only. Let $s(\lambda_k)$ denote the *scheduled* total supply in the k th pricing period, i.e., $s(\lambda_k) = \sum_{i \in S} s_i(\lambda_k)$. We note that, in current electricity wholesale markets, the supply and price are often determined through a bidding process [14], which is generally governed by the costs of generation and transmission. In a competitive bidding-based wholesale market, the resultant supply and price will well reflect the supply model derived from the cost model. We assume that the *realized* total generation in the k th pricing period, denoted by $\hat{s}(\lambda_k)$, is always equal to the total demand $d(k, \lambda_k)$. This is consistent with the current technologies in power grids. For instance, when the demand exceeds the scheduled generation, the system operators will observe a dropping voltage and frequency, and generation can be increased to meet demand and maintain the voltage and frequency at their nominal values.

Empirical supply model: Quadratic cost functions have been widely adopted in the analyses of power generation systems [15], i.e., $c_i(x) = \alpha_i x^2 + \beta_i x + \delta_i$, where $\alpha_i > 0$. To make $c_i(x)$ non-negative over $x \geq 0$, we have a few additional conditions: $\delta_i \geq 0$ if $\beta_i \geq 0$, or $\delta_i \geq \frac{\beta_i^2}{4\alpha_i}$ if $\beta_i < 0$. Therefore, if $\lambda_k \geq \beta_i$, $s_i(\lambda_k) = \frac{1}{2\alpha_i}(\lambda_k - \beta_i)$; otherwise, $s_i(\lambda_k) = 0$. To simplify the evaluation based on this empirical supply model, we assume that $\beta_i < 0$ and $\delta_i \geq \frac{\beta_i^2}{4\alpha_i}$, such that the total supply can be simplified as $s(\lambda_k) = p \cdot \lambda_k + q$, where $p = \sum_{i \in S} \frac{1}{2\alpha_i} > 0$ and $q = \sum_{i \in S} -\frac{\beta_i}{2\alpha_i} > 0$. We now empirically verify this linear supply model using the half-hourly total supply data of New South Wales (NSW), Australia, provided by the Australian Energy Market Operator (AEMO) [1]. Fig. 1 shows the histogram of total supply versus the wholesale price in January, 2012. We can see that the relationship between the average supply and




Figure 2: An example of unstable solution [25] under the linear supply and CEO demand models. Top figure: Evolution of price. Bottom figure: Generation scheduling error in GW. ($p=152, q=4503, \epsilon = -0.6, b_k=2000, \lambda_k^*=20, \lambda_0=21, \lambda_{\min}=1, \lambda_{\max}=100$)

price is nearly linear. A linear fitting of the total supply shown in Fig. 1 yields $p = 152$ and $q = 4503$. Such a linear relationship can also be seen in the investigation of the electricity market of California [29, p. 112], where the demand does not exceed the generation capacity.

4. THE RTP PROBLEM AND SOLUTIONS

This section formally states the RTP problem, examines an existing solution, and proposes a new basic control-theoretic solution with provable *bounded-input bounded-output stability* (referred to as *stability* for short in this paper). Based on our solution, the security analysis in Section 5 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., $s(\lambda_k^*) = d(k, \lambda_k^*)$. However, as $d(k, \lambda_k)$ is unknown, in practice, the ISO sets the price to match the scheduled supply and *predicted* demand (denoted by $\tilde{d}(k, \lambda_k)$). Formally, we define

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

Straightforward solutions to this problem may lead to significantly fluctuating prices. For instance, a direct feedback approach [25], which uses $d(k-1, \lambda_{k-1})$ as the predicted demand $\tilde{d}(k, \lambda_k)$, can yield oscillating prices as shown in Fig. 2. The root cause of the oscillation is the unstable closed-loop system formed by the direct feedback. When the system is unstable, the price set by the ISO will oscillate or diverge, even if the initial price is very close to the true clearing price. The oscillations may lead to severe consequences. 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., $s(\lambda_k) - \hat{s}(\lambda_k)$). Although reserve generating capacity can help compensate for the errors, their use may increase the cost of operating the system.

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. In Section 4.2, we examine the stability of an existing RTP scheme [25] in the absence of attacks. Its poor stability properties motivate us to design a basic control-theoretic RTP scheme with provable stability in the absence of attacks. This is the subject of Section 4.3.



(a) Map of convergence probability ($b = 0$). (b) Stability boundary under various settings of b .

Figure 3: Stability of direct feedback approach under the linear supply and CEO demand models.

4.2 Direct Feedback Approach

A direct feedback approach to the RTP problem has been studied in [25]. The conditions for global stability² of the approach, i.e., the properties of $s(\cdot)$ and $d(\cdot, \cdot)$ that ensure global stability, have also been analyzed in [25]. The approach is briefly reviewed as follows. It predicts $d(k, \lambda_k)$ by the most recent demands based on an autoregression model, and determines the price λ_k accordingly. For instance, the simplest autoregression model uses $d(k-1, \lambda_{k-1})$ as $\hat{d}(k, \lambda_k)$ and the closed-loop system is expressed as $s(\lambda_k) = d(k-1, \lambda_{k-1})$. It is also referred to as the *persistence model*. Thus, the price is determined as $\lambda_k = s^{-1}(d(k-1, \lambda_{k-1}))$. If direct feedback based on the persistence model is not globally stable, it is difficult to stabilize those systems globally with an autoregression-based direct feedback approach [25]. Hence, global stability under the persistence model is particularly important. By applying Corollary 3 in [25], our analysis [30] shows that under the linear supply and CEO demand models, where $\epsilon \in (-1, 0)$ and b_k is a non-negative constant b , the system $\lambda_k = s^{-1}(d(k-1, \lambda_{k-1}))$ is not globally stable.

As the direct feedback approach is not globally stable, its convergence highly depends on the system state. If b_k is time-varying, it can push the system to a state that eventually leads to divergence. A few realistic constraints may affect the system stability. For instance, even if the system is not globally stable, the system may converge when the initial price is within the allowed range $[\lambda_{\min}, \lambda_{\max}]$. Moreover, if a tentative price is out of the range $[\lambda_{\min}, \lambda_{\max}]$, it will be rounded to λ_{\min} or λ_{\max} . Hence, we conduct numerical experiments that account for these realistic constraints for better understanding. The settings of the supply model are $p = 152$ and $q = 4503$, which are obtained in Fig. 1. Given a clearing price $\lambda^* \in [\lambda_{\min}, \lambda_{\max}]$, the coefficient D is set by solving $s(\lambda^*) = d(\lambda^*)$, i.e., $D = \frac{p\lambda^* + q - b}{\lambda^* \epsilon}$, where $b \in [0, p\lambda_{\min} + q]$ to ensure $D > 0$ for any valid λ^* . Fig. 3(a) shows a map of the probability that the system is converging when $b = 0$. To calculate the probability, the initial price sweeps the range $[\lambda_{\min}, \lambda_{\max}]$ and the probability is calculated as the fraction of the initial prices that lead to system convergence. Fig. 3(a) shows that the probability is mostly either 0 or 1 and the transition region with the probability within (0, 1) is sharp. Fig. 3(b) plots the boundaries between the converging and diverging regions under various settings of b , where its valid range is $[0, 4503]$. For instance, when $b = 4000$, the system can be diverging if $\epsilon = -0.8$ and $\lambda^* < 20$. For the data shown in Fig. 1, about 20% of

²The system is *globally stable* if the price converges to the clearing price given any positive initial price.




Figure 4: The control-theoretic price stabilization. $\Lambda(z)$ and $E(z)$ are the z -transforms of λ_k and e_k .

the prices are lower than 20. Therefore, the direct feedback approach can be unstable with significant probabilities.

4.3 Control-Theoretic Price Stabilization

The results in Section 4.2 show the necessity of control laws for stabilizing the RTP systems. This section develops a basic control-theoretic price stabilization algorithm with provable stability. The main objective of this paper is to identify the fundamental impacts of integrity attacks against the vulnerable real-time price signals on the stability of the RTP systems with well designed control laws. More sophisticated price stabilization algorithms could be developed. However, our security analysis in Section 5, which is based on our basic control-theoretic algorithm, provides fundamental baselines for understanding the security properties of RTP systems running such sophisticated algorithms.

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 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. 4. We let $G_c(z)$, $G_p(z)$, and $H(z)$ denote the transfer functions of the price stabilization algorithm, the controlled system, and the observation system, which are expressed in the z -transform domain. The z -transform [24] 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 modelled as a *disturbance* to the system [24].

We now derive the expressions of $G_p(z)$ and $H(z)$. To preserve generality, our design is based on the abstract supply and demand models $s(\lambda)$ and $w(\lambda)$, which can be non-linear as in the CEO model. In controller design, a common approach to dealing with non-linear systems is to adopt local linearization [24]. 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. 4, 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 price stabilization 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. Based on the above modeling, we have the following proposition, which can be proved by examining whether the poles of the system are located within a unit circle centered at the origin of the z -plane. The details of the proof are omitted due to space constraints and can be found in [30].

PROPOSITION 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 price stabilization 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 transfer function of the above algorithm is $G_c(z) = \frac{2\eta}{(\dot{s}(\lambda_o) - \dot{w}(\lambda_o))(1 - z^{-1})}$. From control theory, when b_k is a constant, the system converges the fastest when $\eta = 0.5$, as the system's pole is at the origin [24]. The convergence speed is particularly 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 Section 5 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 in Section 6.1 shows that, if the relative error in estimating $\dot{w}(\lambda_o)$ is less than $100 \times (1 - \eta)\%$, the algorithm given by Proposition 1 remains stable. 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 1, we do not elaborate on the estimation algorithm, and the analysis in Section 5 assumes that the ISO can accurately estimate $\dot{w}(\lambda_o)$. Section 6.1 also discusses the impact of inaccurate $\dot{w}(\lambda_o)$ on the security analysis.

The price stabilization algorithm in Proposition 1 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 $w(\lambda)$ can be more accurate. Specifically, by setting $\lambda_o = \lambda_{k-1}$, we have the following algorithm:

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

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 the settings shown in Fig. 3. The numerical examples and simulations conducted in the rest of this paper employ the algorithm in Eq. (2). Fig. 5(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 but converges. The oscillation is caused by a negative pole [24]. Fig. 5 will also be used as a running example in Section 5 to illustrate the impact of attacks.

5. INTEGRITY ATTACKS TO RTP

This section studies the impact of two integrity attacks on RTP systems under the RTP scheme given by Proposition 1.




Figure 5: A running example under the linear supply and CEO demand models ($T = 0.5$, $\epsilon = -0.8$, $b_k = 2000$, $\lambda^* = 20$, $\lambda_0 = 21$). (a) Price stabilization; (b) Scaling attack ($\eta = 0.8$, $\rho = 1$, $\gamma = 0.57$); (c) Delay attack ($\eta = 0.2$, $\rho = 1$, $\tau = 12$).

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 smart 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 [22, 26] 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. In this paper, as the price signals sent to the centralized suppliers are often well protected, we assume that they are not subject to attacks. However, our analysis framework can be easily extended to account for possible attacks on the suppliers.

5.1.1 Attack Models

As the number of consumers in a smart 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

$$\rho = \frac{\sum_{j \in C'} w_j(\lambda'_k)}{\sum_{j \in C} w_j(\lambda'_k)} = \frac{\sum_{j \in C'} w_j(\lambda'_k)}{w(\lambda'_k)}, \quad (3)$$

which characterizes the fraction of consumers receiving the compromised price signals. If the consumers are *homoge-*

neous (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 [30] shows that, if the heterogeneous consumers follow the CEO model, $\rho \simeq |C'|/|C|$ with a variation of less than 0.003 and hence can be practically treated as a constant. Moreover, we make the following approximation:

$$\sum_{j \in C \setminus C'} w_j(\lambda_k) \simeq (1-\rho) \sum_{j \in C} w_j(\lambda_k) = (1-\rho)w(\lambda_k). \quad (4)$$

The numerical evaluation in [30] shows that relative approximation error of Eq. (4) is less than 1%. Therefore, in the presence of integrity attacks, we have

$$w'(\lambda_k) \simeq \rho w(\lambda'_k) + (1-\rho)w(\lambda_k). \quad (5)$$

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:

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

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 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 [28] synchronize their clocks via a built-in GPS receiver, which is vulnerable and subject to remote attacks that are effective across large geographic areas [23].

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 6.2, 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 1. *Marginal demand-supply ratio is $h = \left| \frac{\dot{w}(\lambda_o)}{\dot{s}(\lambda_o)} \right|$.*

From Definition 1, h depends on the operating point λ_o . As discussed in Section 4.3, the gain coefficient η of the price stabilization algorithm affects the system stability in a major way. Therefore, we define the following metric:

DEFINITION 2. *Given attack \mathcal{A} , the region of operating point stability under attack, denoted by $\text{ROS}_{\lambda_o}(\mathcal{A})$, is*

$$\text{ROS}_{\lambda_o}(\mathcal{A}) = \{(h, \eta) | \text{The system is stable under attack } \mathcal{A}\}.$$

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

DEFINITION 3. *Given attack \mathcal{A} , the region of stability under attack, denoted by $\text{ROS}(\mathcal{A})$, is*

$$\text{ROS}(\mathcal{A}) = \{\eta | \text{The system is stable under attack } \mathcal{A}, \forall h > 0\}.$$

The above two metrics are important for understanding the impact of integrity attacks on the stability of the RTP system under the price stabilization algorithm in Proposition 1. In particular, the $\text{ROS}(\mathcal{A})$ 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 $\text{ROS}(\mathcal{A}_1) \subset \text{ROS}(\mathcal{A}_2)$, 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 $\eta \in \text{ROS}(\mathcal{A})$ 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 Attack

The local linearization of Eq. (5) with $\lambda'_k = \gamma \lambda_k$ is

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

By collecting the price-independent terms with b_k , the transfer function of the price-dependent component is $G_w(z) = \rho \gamma \dot{w}(\gamma \lambda_o) + (1-\rho) \dot{w}(\lambda_o)$. To make the analysis tractable, for the scaling attack, we only focus on a class of price-responsive demand models that satisfy $\dot{w}(\gamma \lambda) = \dot{w}(\lambda) \cdot \mu(\gamma | \Theta)$, where Θ is the set of model parameters of $w(\cdot)$ and the function $\mu(\gamma | \Theta)$ is independent of λ and always positive. Such a $\dot{w}(\cdot)$ is said to be *decomposable*. For instance, under the CEO model, $\Theta = \{D, \epsilon\}$, and $\mu(\gamma | \Theta) = \gamma^{\epsilon-1}$. For simplicity of exposition, we denote $\mu(\gamma | \Theta)$ as μ in the rest of this paper. Therefore, $G_w(z) = \rho \gamma \mu \dot{w}(\lambda_o) + (1-\rho) \dot{w}(\lambda_o)$, and $G_p(z) = G_s(z) - G_w(z) = \dot{s}(\lambda_o) - \rho \gamma \mu \dot{w}(\lambda_o) - (1-\rho) \dot{w}(\lambda_o)$. The closed-loop transfer function [24] under the attack is

$$T_c(z) = \frac{G_c(z)G_p(z)}{1+G_c(z)G_p(z)H(z)} = \frac{2\eta(1+\rho\gamma\mu h+h-\rho h)z}{P(z)},$$

where the *system characteristic function* $P(z) = (h+1)(z-1) + 2\eta(1+\rho\gamma\mu h+h-\rho h)$. Note that $G_p(z)$, $H(z)$, and $G_c(z)$ have been obtained in Section 4.3.




Figure 6: Stability boundaries under scaling attack, abstract supply model, and CEO demand model ($\epsilon = -0.8$, ROS_{λ_o} are the regions below the boundaries).

5.2.1 Region of Operating Point Stability

PROPOSITION 2. *For the linearized system based on a fixed operating point λ_o and a decomposable $\dot{w}(\cdot)$, $\text{ROS}_{\lambda_o}(\rho, \gamma) = \{(h, \eta) | 0 < \eta < \min\{1, \bar{\eta}\}, \forall h > 0\}$, where*

$$\bar{\eta} = \frac{h+1}{h+1+\rho h(\gamma\mu-1)}. \quad (6)$$

PROOF. If all the poles of $T_c(z)$ (i.e., roots of $P(z)$) are within the unit circle centered at the origin of z -plane, the system is stable [24]. If $\eta < \bar{\eta}$, the pole is within the circle. As $\eta \in (0, 1)$, η takes the minimum of 1 and $\bar{\eta}$. \square

Remark 1. Under the CEO demand model, by replacing $\mu = \gamma^{\epsilon-1}$ in Eq. (6), we have $\bar{\eta} = \frac{h+1}{h+1+\rho h(\gamma^{\epsilon-1}-1)}$. Fig. 6 plots the stability boundaries when $\epsilon = -0.8$, 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 now use the numerical example in Fig. 5(b) to verify our analysis. Fig. 5(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 analytical ROS_{λ_o} . Therefore, Proposition 2 successfully characterizes the critical stability boundary. Note that, as the settings for Fig. 5(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 7.

5.2.2 Region of Stability

PROPOSITION 3. *For the linearized system based on a decomposable $\dot{w}(\cdot)$, when $\gamma\mu \in (0, 1]$, $\text{ROS}(\rho, \gamma) = \{\eta | 0 < \eta < 1\}$; when $\gamma\mu > 1$, $\text{ROS}(\rho, \gamma) = \{\eta | 0 < \eta < \inf_{h>0} \bar{\eta}\}$, where $\inf_{h>0} \bar{\eta} = \frac{1}{1+\rho(\gamma\mu-1)}$.*

PROOF. When $\gamma\mu \in (0, 1]$, $\bar{\eta} \geq 1$. From Proposition 2, if $0 < \eta < 1$, the system is stable regardless of h . When $\gamma\mu > 1$, $\bar{\eta}$ is a bounded decreasing function of h . Its infimum $\inf_{h>0} \bar{\eta} = \lim_{h \rightarrow +\infty} \bar{\eta} = \frac{1}{1+\rho(\gamma\mu-1)}$. Therefore, if $0 < \eta < \inf_{h>0} \bar{\eta}$, the system is stable regardless of h . \square




Figure 7: Stability boundaries under delay attack and abstract supply/demand models. (In the left figure, the curve with $\tau = 1$ and $\gamma = 0.6$ is for combined attack discussed in Section 6.2.)

Remark 2. Under the CEO demand model, replacing $\mu = \gamma^{\epsilon-1}$ in Proposition 3 yields the following result. When $\gamma \geq 1$, $\text{ROS}(\rho, \gamma) = \{\eta | 0 < \eta < 1\}$; when $\gamma \in (0, 1)$, $\text{ROS}(\rho, \gamma) = \{\eta | 0 < \eta < \inf_{h>0} \bar{\eta}\}$, where $\inf_{h>0} \bar{\eta} = \frac{1}{1+\rho(\gamma^{\epsilon-1}-1)}$. Therefore, under the CEO model, 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(a) plots $\inf_{h>0} \bar{\eta}$ when $\epsilon = -0.8$. We can see that ROS shrinks with increased ρ and decreased γ . This can be easily proved by the monotonicity of $\inf_{h>0} \bar{\eta}$.

5.3 Impact of Delay Attack

The local linearization of Eq. (5) 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 $z^{-\tau}$ represent 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 [24]. Among them, the Jury test [24, 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. 7 plots the stability boundaries, where the ROS_{λ_o} are the regions below the boundaries. From Fig. 7, 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 is omitted due to space constraints and can be found in [30].

Algorithm 1 Compute $\text{ROS}(\rho, \tau)$ when $\rho \in (0.5, 1]$

Input: ρ and τ
Output: $\lim_{h \rightarrow +\infty} \bar{\eta}(h|\rho, \tau)$
1: **if** $\tau = 1$ **then**
2: **return** $1/(2 \cdot \rho)$
3: **end if**
4: $\mathbf{X} = 1, \mathbf{Y} = 2\eta\rho, \mathbf{Z} = 2\eta\rho(1 - 2\eta(1 - \rho))$
5: **for** $i = 1$ **to** τ **do**
6: $\mathbf{U} = \mathbf{X} \cdot \mathbf{X} - \mathbf{Y} \cdot \mathbf{Y}, \mathbf{V} = \mathbf{Z}, \mathbf{W} = (\mathbf{X} \cdot \mathbf{Z} \cdot \mathbf{Z})/\mathbf{Y}$
7: $\mathbf{X} = \mathbf{U}, \mathbf{Y} = \mathbf{V}, \mathbf{Z} = \mathbf{W}$
8: **end for**
9: $\mathbf{Q}(\eta|\rho) = \mathbf{X} \cdot \mathbf{X} - \mathbf{Y} \cdot \mathbf{Y} - \mathbf{Z}$
10: **return** minimum root of $\mathbf{Q}(\eta|\rho) = 0$ over $\eta \in (0, 1)$

Note: Line 4 to Line 9 are symbolic calculation, where the bold capitals are symbolic expressions of η and ρ .

PROPOSITION 4. For the linearized system with a fixed operating point λ_o , $\text{ROS}_{\lambda_o}(\rho, \tau + 1) \subseteq \text{ROS}_{\lambda_o}(\rho, \tau)$.

We now use the numerical example in Fig. 5 to verify our analysis. Fig. 5(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. Note that when $\tau = 11$, the price does converge and the average value of h is 1.455335, which falls in the stable region ($h < 1.522$) according to the Jury test. Therefore, the Jury test successfully characterizes the critical stability boundary. As the settings for Fig. 5(c) are close to the stability boundary, the price diverges slowly. For larger τ , the price can diverge quickly.

5.3.2 Region of Stability

We observe from Fig. 7(b) that, when $\rho = 0.5$, the system is stable for $\eta \in (0, 1)$. We have the following proposition.


PROPOSITION 5. For the linearized system, if $\rho \in (0, 0.5]$, $\forall \tau \in \mathbb{Z}^+$, $\text{ROS}(\rho, \tau) = \{\eta | 0 < \eta < 1\}$.

The proof can be found in the appendix, where we prove that if $\rho \in (0, 0.5]$, all roots of $P(z)$ are within the unit circle centered at the origin in the z -plane and hence the system is stable [24]. From Proposition 5, to launch a successful delay attack that destabilizes the system, the adversary has to compromise no less than a 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 the 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 price stabilization algorithm given by Proposition 1 to delay attacks.

We now discuss the ROS when $\rho \in (0.5, 1]$. From Fig. 7, the stability boundary curves are non-increasing and converge to limits when $h \rightarrow +\infty$. Let $\bar{\eta}(h|\rho, \tau)$ denote the stability boundary curve for particular ρ and τ . Therefore,

$$\text{ROS}(\rho, \tau) = \{\eta | 0 < \eta < \lim_{h \rightarrow +\infty} \bar{\eta}(h|\rho, \tau)\}.$$

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. Instead, we use an algorithm to define $\text{ROS}(\rho, \tau)$,



(a) Scaling attack (abstract supply model and CEO demand model with $\epsilon=-0.8$) (b) Delay attack (abstract supply and demand models)

Figure 8: Upper bound of ROS.

which is shown in Algorithm 1. This algorithm is developed based on key observations from the Jury test procedure. Fig. 8(b) plots $\lim_{h \rightarrow +\infty} \bar{\eta}(h|\rho, \tau)$, which is computed by Algorithm 1, versus τ under various settings of ρ . We also use the Jury test to compute $\bar{\eta}(h|\rho, \tau)$ with a large setting for h (specifically, $h = 10^{10}$). The results are the same as in Fig. 8(b). From the figure, we can see that the ROS shrinks with ρ and τ , which is consistent with intuition.

6. DISCUSSIONS

In this section, we discuss the impact of inaccuracy in estimating $\dot{w}(\lambda_o)$ on the analysis in the previous sections. We also discuss how to extend our analysis to address more complicated attack models.

6.1 Impact of Inaccuracy in Estimating $\dot{w}(\lambda_o)$

As the price-dependent demand model $w(\cdot)$ is unknown to the ISO, we derive an upper bound for the error in estimating $\dot{w}(\lambda_o)$, to ensure the stability of the algorithm in Proposition 1. Let $\widetilde{\dot{w}}(\lambda_o)$ denote the estimated $\dot{w}(\lambda_o)$, and $E_w = \frac{\dot{w}(\lambda_o) - \widetilde{\dot{w}}(\lambda_o)}{\dot{w}(\lambda_o)}$ denote the relative estimation error. The stability condition $0 < \eta < 1$ can be rewritten as $0 < \frac{2\eta}{\dot{s}(\lambda_o) - \dot{w}(\lambda_o)} < \frac{2}{\dot{s}(\lambda_o) - \dot{w}(\lambda_o)}$. As long as $0 < \frac{2\eta}{\dot{s}(\lambda_o) - \dot{w}(\lambda_o)} < \frac{2}{\dot{s}(\lambda_o) - \dot{w}(\lambda_o)}$, the system is stable. This condition can be derived as $E_w < (1-\eta) \left(1 - \frac{\dot{s}(\lambda_o)}{\dot{w}(\lambda_o)}\right)$. As $1 - \frac{\dot{s}(\lambda_o)}{\dot{w}(\lambda_o)} > 1$, $E_w < (1-\eta)$ is a sufficient condition for stability.

We now discuss the impact of inaccurate $\widetilde{\dot{w}}(\lambda_o)$ on the security analysis results in Section 5. From the definition of E_w , we have $\widetilde{\dot{w}}(\lambda_o) = (1 - E_w)\dot{w}(\lambda_o)$. Note that $E_w < 1$ since $\dot{w}(\lambda_o) > 0$. By replacing h with $|1 - E_w| \cdot h$ in Proposition 2, we have a new result in the presence of the estimation error E_w . From the proofs of Propositions 3, 4, and 5, they are independent of h . Therefore, these propositions still hold in the presence of estimation errors.

6.2 Superimposed and Heterogeneous 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.

From discrete-time control theory [24], 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 studied in Section 5 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, the closed-loop system characteristic function of Eq. (5) 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 $\text{ROS}_{\lambda_o}(\rho, \gamma, \tau)$ and $\text{ROS}(\rho, \gamma, \tau)$. Fig. 7(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 in the grid, Eq. (5) 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.

7. TRACE-DRIVEN SIMULATIONS

We use GridLAB-D [2], an electric power distribution system simulator, to evaluate the impact of integrity attacks. GridLAB-D captures many physical characteristics such as power line capacities and impedances. Hence, we can validate our analysis under the realism provided by GridLAB-D. Moreover, it can record emergency events that occur when the current ratings of lines and power ratings of transformers are exceeded. Such events could cause sustained service interruptions to consumers. These events are of particular interest to us, because they help us understand the physical consequences caused by the integrity attacks.

7.1 Simulation Methodology and Settings

We use a distribution feeder specification [27], which covers a moderately populated urban area and comprises 1405 houses, 2134 buses, 3314 triplex buses, 1944 transformers, 1543 overhead lines, 335 underground lines, and 1631 triplex lines. For this small-scale distribution feeder, locational prices 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 price stabilization 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




Figure 9: Price stabilization in the absence of attack.

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: $D_i \sim \mathcal{N}(7, 3.5^2)$ (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. To improve the realism of the simulations, we use the half-hourly total demand trace from March 1st to 22nd, 2013, of NSW, Australia, provided by AEMO [1] as the baseline load. The baseline load of a single house is set to be a scaled version of the real data trace. 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 [3]. In our simulations, the price is updated every half an hour, to be consistent with the setting of the demand data traces [1]. 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 1405 houses. From the facts that there are 2.8 million households in NSW [3] and 57% of AEMO's supply is for residential demand [9, 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 $\eta = 0.5$.

7.2 Simulation Results

Price stabilization: The first set of simulations evaluate the effectiveness of the direct feedback approach [25] and our control-theoretic price stabilization algorithm in Eq. (2). In the simulations, the direct feedback approach is unstable, where the price oscillates between λ_{\min} and λ_{\max} . The total demand reaches 10 MW a few hours after the start of the simulation, and GridLAB-D reports that four distribution lines are overloaded. Fig. 9 plots the price and resultant demands under our price stabilization approach 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. We also evaluate the impact of an inaccurate consumer model $w(\lambda) = \tilde{D}\lambda^{\tilde{\epsilon}}$ on the system performance. The evaluation results are consistent with our analysis in Section 6.1. For instance, when $\eta = 0.5$, the system keeps stable if i) $(D - \tilde{D})/D < 0.67$ and $\tilde{\epsilon} = \epsilon$, or ii) $0.05\epsilon < \tilde{\epsilon} < 2.15\epsilon$ and $\tilde{D} = D$. Therefore, the price stabilization algorithm is resilient to inaccuracies of the estimated consumer model.

Scaling attack: Fig. 10 plots the true and compromised prices, as well as the breakdown of demand under the scal-




Figure 10: Scaling attack ($\rho = 65\%$, $\gamma = 0.1$).





Figure 11: Impact of delay attack ($\rho = 100\%$, $\tau = 9$). (a) Prices and demands; (b) emergencies.

ing attack. We can see that the price as well as the demand fluctuates severely. GridLAB-D reports excessive distribution 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 metric. A near-zero $\sigma(e)$ means convergence, while a considerably large $\sigma(e)$ means oscillation or divergence. Fig. 13(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 .

Delay attack: Fig. 11(a) and Fig. 12(a) show the evolution of price and the breakdown of demand under the delay attacks with different parameters. We also investigate the emergency events reported by GridLAB-D. The *overload* of a distribution line or a transformer is defined as the percentage of the exceeded current/power with respect to the rated value. Fig. 11(b) and Fig. 12(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. 11, 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. 12, 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. Lastly, we evaluate the impact of the delay attack with different settings of ρ and τ . The results are shown in Fig. 13(b). We can clearly see that when $\rho < 0.5$, the system remains stable, which is consistent with Proposition 5.

8. CONCLUSION

This paper systematically investigates the impact of scaling and delay attacks on the stability of RTP systems. We






Figure 12: Impact of delay attack ($\rho = 65\%$, $\tau = 24$). (a) Prices and demands; (b) emergencies.



(a) Scaling attack (b) Delay attack

Figure 13: System volatility under attacks.

characterize the impact using a control-theoretic metric, *region of stability*. We show that, to destabilize the RTP system, it is necessary for the adversary to reduce the price in a scaling attack or compromise more than half of the smart meters 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.

8.1 Limitations and Future Work

In this paper, we made a few simplifying assumptions as we pointed out. Care is thus required when applying our results to general real-world settings, e.g., large-scale smart grids, where locational prices are important. Particularly, Propositions 1 to 5 are based on the linearized abstract supply and demand models at a fixed operating point.³ The security analysis in Propositions 2 to 5 is based on the approximate aggregate demand model in Eq. (5). These simplifications enable us to study the RTP susceptibility to malicious attacks under LTI settings. However, for systems that involve non-linear, time-variant, or stochastic components and complex market structures, the LTI assumptions may not hold and our analysis may lead to inaccuracies in characterizing such systems.

It is interesting for further research to address the following issues not considered in this paper. First, our analysis framework can be extended to address attacks that cannot be modeled as LTI systems. Second, for large-scale systems where locational prices are important, the system behavior may deviate from our analysis when the power network is congested due to price oscillations. Locational prices can be considered by integrating economic dispatch into the analysis. Third, this paper considers simplified system models while preserving the principles of RTP. Refinements are possible that consider various practical factors such as suppliers' ramp constraints, energy storage, load shifting, bidding markets, and ex-ante RTP with ex-post adjustments.

³The numerical results and simulations use the algorithm in Eq. (2), which uses the current price as the operating point.

9. ACKNOWLEDGEMENTS

We thank our shepherd Dr. Alvaro A. Cárdenas and the anonymous reviewers for providing valuable feedback on this work. This work was supported in part by Singapore's Agency for Science, Technology and Research (A*STAR) under the Human Sixth Sense Programme, and in part by U.S. National Science Foundation under grant numbers CNS-0963715 and CNS-0964086.

10. REFERENCES

- [1] Australian energy market operator. www.aemo.com.au.
- [2] GridLAB-D. www.gridlabd.org.
- [3] How energy is used in the home. www.bit.ly/h9YHzV.
- [4] Multiple vulnerabilities in NTP. www.bit.ly/GzDJ5a.
- [5] PJM operational data. www.bit.ly/diyZ9t.
- [6] RTP-HA-2 program. www.georgiapower.com.
- [7] F. Alvarado. The stability of power system markets. *IEEE Trans. Power Syst.*, 14(2):505–511, 1999.
- [8] S. Amin, X. Litrico, S. Sastry, and A. M. Bayen. Cyber security of water SCADA systems-Part I: Analysis and experimentation of stealthy deception attacks. *IEEE Trans. Control Syst. Technol.*, 2012.
- [9] Australian Energy market operator. 2011 national electricity forecasting. www.bit.ly/ZQzk0s.
- [10] G. Barbose, C. Goldman, R. Bharvirkar, N. Hopper, M. Ting, and B. Neenan. 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 Laboratory, 2005.
- [11] A. A. Cárdenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry. Attacks against process control systems: risk assessment, detection, and response. In *ASIACCS*, 2011.
- [12] A. A. Cárdenas, S. Amin, and S. Sastry. Secure control: Towards survivable cyber-physical systems. In *1st Intl. Workshop on Cyber-Physical Systems*, 2008.
- [13] M. Filippini. Short-and long-run time-of-use price elasticities in Swiss residential electricity demand. *Energy Policy*, 39(10):5811–5817, 2011.
- [14] S. Fleten and E. Pettersen. Constructing bidding curves for a price-taking retailer in the norwegian electricity market. *IEEE Trans. Power Syst.*, 20(2):701–708, 2005.
- [15] J. Grainger and W. Stevenson. *Power system analysis*. McGraw-Hill, 1994.
- [16] L. Jia, R. J. Thomas, and L. Tong. Impacts of malicious data on real-time price of electricity market operations. In *Intl. Conf. System Sciences*, 2012.
- [17] O. Kosut, L. Jia, R. J. Thomas, and L. Tong. Malicious data attacks on the smart grid. *IEEE Trans. Smart Grid*, 2(4):645–658, 2011.
- [18] P. Kundur. *Power system stability and control*. Tata McGraw-Hill Education, 1994.
- [19] M. G. Lijesen. The real-time price elasticity of electricity. *Energy economics*, 29(2):249–258, 2007.
- [20] J. Lin, W. Yu, X. Yang, G. Xu, and W. Zhao. On false data injection attacks against distributed energy routing in smart grid. In *ICCPS*, 2012.
- [21] Y. Liu, P. Ning, and M. K. Reiter. False data injection attacks against state estimation in electric power grids. In *CCS*, 2009.

- [22] S. McLaughlin, D. Podkuiko, S. Miadzvezhanka, A. Delozier, and P. McDaniel. Multi-vendor penetration testing in the advanced metering infrastructure. In *ACSAC*, 2010.
- [23] T. Nighswander, B. Ledvina, J. Diamond, R. Brumley, and D. Brumley. GPS software attacks. In *CCS*, 2012.
- [24] K. Ogata. *Discrete-time control systems*. Prentice-Hall Englewood Cliffs, NJ, 2 edition, 1995.
- [25] M. Roozbehani, M. Dahleh, and S. Mitter. Volatility of power grids under real-time pricing. *IEEE Trans. Power Syst.*, 27(4):1926–1940, 2012.
- [26] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser. Neighborhood watch: security and privacy analysis of automatic meter reading systems. In *CCS*, 2012.
- [27] K. P. Schneider, Y. Chen, D. P. Chassin, R. Pratt, D. Engel, and S. Thompson. Modern grid initiative distribution taxonomy final report. Technical report, Pacific Northwest National Laboratory, 2008.
- [28] Schneider Electric. Technical note of ION smart meter. www.bit.ly/15W9GR4.
- [29] J. Sweeny. *The California electricity crisis*. Hoover Institution Press, 2002.
- [30] R. Tan, V. Badrinath Krishna, D. K. Y. Yau, and Z. Kalbarczyk. Impact of integrity attacks on real-time pricing in smart grids. Technical report, Advanced Digital Sciences Center, 2013. <http://publish.illinois.edu/cps-security/files/2013/08/stability.pdf>.
- [31] The Wall Street Journal. Spies penetrate U.S. electric grid. <http://on.wsj.com/fhGdG>.
- [32] L. Xie, Y. Mo, and B. Sinopoli. Integrity data attacks in power market operations. *IEEE Trans. Smart Grid*, 2(4):659–666, 2011.
- [33] Y. Yuan, Z. Li, and K. Ren. Modeling load redistribution attacks in power systems. *IEEE Trans. Smart Grid*, 2(2):382–390, 2011.

APPENDIX: PROOF FOR PROPOSITION 5

Denote $u_1 = h+1$, $u_2 = 2\eta + 2\eta(1-\rho)h - h - 1$, $u_3 = 2\eta\rho h$, and $P(z) = u_1 z^{\tau+1} + u_2 z^\tau + u_3$, where $u_1 > 0$ and $u_3 > 0$. Express any system pole (i.e., root of $P(z) = 0$) in polar coordinates: $z = A(\cos\theta + i\sin\theta)$, where $A > 0$. $P(z) = 0$ can be rewritten as two equations: $A^\tau(u_1 A \cos(\tau+1)\theta + u_2 \cos\tau\theta) = -u_3$ and $A^\tau(u_1 A \sin(\tau+1)\theta + u_2 \sin\tau\theta) = 0$. Adding the squares of the two equations yields $g(A) = 0$, where $g(A) = u_1^2 A^{2\tau+2} + 2u_1 u_2 \cos\theta A^{2\tau+1} + u_2^2 A^{2\tau} - u_3^2$. Thus, any pole satisfies $g(A) = 0$. We can verify that $g(1) > 0$ when $\rho \in (0, 0.5]$. Moreover, $\dot{g}(A) = A^{2\tau} m(A)$, where $m(A) = u_1^2(2\tau+2)A + 2u_1 u_2(2\tau+1)\cos\theta + \frac{2\tau u_2^2}{A}$ is a convex function with its minimum at $A^* = \left| \frac{u_2}{u_1} \right| \sqrt{\frac{\tau}{\tau+1}}$. We can verify that $\left| \frac{u_2}{u_1} \right| < 1$ if $\rho \in (0, 0.5]$. Thus, $A^* < 1$ and the minimum of $m(A)$ for $A \geq 1$ is $m(1)$, which satisfies $m(1) \geq u_1^2(2\tau+2) - 2u_1|u_2|(2\tau+1) + 2\tau u_2^2 = 2(u_1 - |u_2|)(\tau(u_1 - |u_2|) + u_1)$. If $u_2 < 0$, $u_1 - |u_2| = 2\eta + 2\eta(1-\rho)h > 0$; if $u_2 \geq 0$, $u_1 - |u_2| = 2(1-\eta) + 2h(1-\eta(1-\rho)) > 0$. Hence, $m(A) > 0$ and $\dot{g}(A) > 0$ for $A \geq 1$. Recalling $g(1) > 0$, we have $g(A) > 0$ for $A \geq 1$. Hence, $A < 1$ for all poles and the system is stable.