

Cyber-Attacks and Mitigation in Blockchain Based Transactive Energy Systems

Carlos Barreto*, Taha Eghtesad[†], Scott Eisele*, Aron Laszka[†], Abhishek Dubey*, Xenofon Koutsoukos*

*Vanderbilt University · {carlos.a.barreto, scott.r.eisele, abhishek.dubey, xenofon.koutsoukos}@vanderbilt.edu

[†]University of Houston · {teghtesad, alaszka}@uh.edu

Published in the proceedings of the
3rd IEEE International Conference on
Industrial Cyber-Physical Systems (ICPS 2020)

Abstract—Power grids are undergoing major changes due to the rapid adoption of intermittent renewable energy resources and the increased availability of energy storage devices. These trends drive smart-grid operators to envision a future where peer-to-peer energy trading occurs within microgrids, leading to the development of Transactive Energy Systems. Blockchains have garnered significant interest from both academia and industry for their potential application in decentralized TES, in large part due to their high level of resilience. In this paper, we introduce a novel class of attacks against blockchain based TES, which target the gateways that connect market participants to the system. We introduce a general model of blockchain based TES and study multiple threat models and attack strategies. We also demonstrate the impact of these attacks using a testbed based on GridLAB-D and a private Ethereum network. Finally, we study how to mitigate these attack.

Index Terms—Transactive Energy System, Blockchain, Cyber-security, Cyber-Physical System, Denial of Service, Microgrid

I. INTRODUCTION

Power grids are undergoing major changes due to the rapid adoption of intermittent renewable resources (e.g., wind and solar), combined with energy storage devices (e.g., residential batteries and electric vehicles). Also, Internet of Things (IoT) devices enable better management of loads and energy resources. These trends augment the capabilities of residential users, now called *prosumers*, because they can both produce and consume energy.

Smart grid operators envision a future where prosumers trade energy or services without intermediaries, improving the efficiency and reliability of power systems. Thus, future grids will use Transactive Energy Systems (TES) as a distributed management approach in which smart appliances or Internet of things (IoT) devices participate autonomously in electricity markets [1]. TES allows smart appliances to assess the energy prices in order to adjust their load reducing costs. Likewise, prosumers can trade their surplus energy with neighbors.

These TES can use either centralized or decentralized markets. In a centralized market, all participants communicate with a central entity, which collects bids and returns the energy price (and the transactions among participants). Centralized markets suffer from a *single point of failure*, because they rely on a single trusted entity to operate the market. Decentralized markets based on blockchains offer several desirable proper-

ties in energy applications. First, prosumers interact without intermediaries and conflicts are resolved through protocols. Second, transactions that have been recorded on the blockchain are immutable and publicly auditable by design. Third, the blockchain is fault tolerant, that is, it can operate even if some of the prosumers fail or act maliciously. These properties can ensure market transparency, as well as the availability of detailed information about the system.

More recent blockchain implementations, such as Ethereum [2], also enable trustworthy computations through *smart contracts* [3]. Based on this functionality, these blockchains can implement various data verification and market clearing mechanisms for TES [4]. For example, smart contracts can enforce commitments as well as transfer of assets between peers.

One benefit of blockchain based TES is resilience: to disrupt the integrity of the market (e.g., tamper with bids or with the clearing mechanism), an attacker needs to compromise a large number of blockchain nodes. A blockchain based system can also resist availability attacks, since the market remains operational even with many unavailable nodes [5]. However, some attacks may degrade the operation of the system.

In practice, IoT devices lack resources required for participating in the computing-intensive consensus algorithms of many blockchains. Thus, prosumers have to connect to a blockchain-based system through *gateway* nodes; however, an adversary can attempt to “cut off” prosumers from the system by targeting these gateway nodes. For example, an adversary can launch a (*distributed*) *denial of service* (DDoS) attack against a gateway node to prevent a set of bids from arriving at the market, which change the market’s equilibria.

In this paper, we study *blockchain based Transactive Energy Systems* and introduce a *novel class of attacks* that target the gateways between prosumers and the system. The following are our main contributions:

- We formulate a general model of blockchain based transactive energy systems, which includes both infrastructure and market mechanisms.
- We introduce a previously unconsidered class of attacks, which *discard or delay trading bids*. Our threat model includes three scenarios, which consider distinct capabilities and knowledge for the adversary.
- We study attack strategies for each scenario. We also discuss how to *mitigate* these attacks by taking advantage of the distributed nature of the system.

- Finally, we present our testbed based on GridLab-D [6], a power system simulator, and a private Ethereum network [7]. We show that attacks on miners can change the market equilibria benefiting adverse generators, but they can be mitigated using the proposed approach.

The remainder of this paper is organized as follows. In Section II, we discuss related work on cyber-attacks against energy systems. In Section III, we introduce our system model, which includes the infrastructure as well as the market mechanisms. In Section IV, we introduce our threat model, specifying the adversary’s capabilities and goal. In Section V, we investigate adversarial strategies and discuss mitigation. In Section VI, we present our testbed and experimental results on attacks and mitigation. Finally, in Section VII, we provide concluding remarks.

II. RELATED WORK

Recent cyber attacks against critical infrastructure, such as the attacks on the Ukrainian power grid in 2015 and 2016 [8], have motivated multiple research efforts to protect critical infrastructures, in particular, the power grid [9].

Prior works have shown how *false data injection* (FDI) attacks can modify sensor measurements to induce errors in a power system’s operation [10], [11]. With a careful design, these attacks can damage the system or change the electricity prices. An adversary can also affect forecast systems, which are used to plan the power-system operation, by exploiting vulnerabilities of artificial intelligence models [12], [13].

In most cases, FDI attacks need information about the state of the system or the models used for making decisions (e.g., the system’s state, its topology, or prediction models). However, some attacks leverage the market’s infrastructure to bypass these restrictions. For example, an adversary that compromises bids can induce changes in the market’s equilibria without knowing details of the system [14], [15].

DDoS attacks represent a significant threat for distributed electricity markets, because an adversary needs minimal knowledge (and resources) to mount attacks. Furthermore, with these attacks, it is extremely difficult to determine the identity of the adversary. For example, [16] reported that a company specializing in protection against DDoS attacks co-authored the Mirai malware to attack some of its customers.

New technologies, such as Internet of Things (IoT) devices, introduce vulnerabilities for the power grid [17], [18]. As a result, adversaries can target customer-side components, such as smart meters, appliances, end-user generation systems (e.g., solar panels), and electric vehicles, to affect the power system’s operation [19]. For example, adversaries can compromise IoT devices to change their bids [15].

Our work is in line with these efforts to improve the protection of critical infrastructure. Specifically, we analyze how blockchain based markets, which are often presented as a resilient solution, can still be attacked by exploiting market mechanisms.

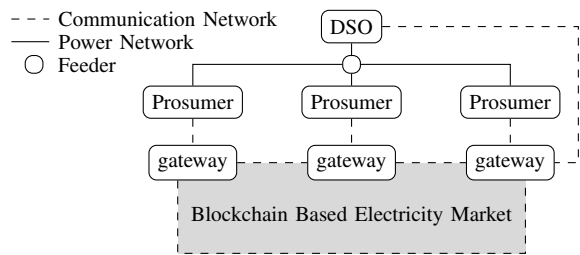


Fig. 1: Decentralized Transactive Energy System infrastructure.

III. SYSTEM MODEL

In this section, we present our system model for a decentralized TES. We make some assumptions based on the inspection of various industrial implementations of decentralized TES, such as LO3 [20] and Power Ledger [21], and scientific articles, such as Laszka *et al.* [22], [23] and Wörner *et al.* [24].

A. Infrastructure

Fig. 1 shows the overall architecture of the decentralized TES. Below we describe each component.

1) *Prosumers*: Agents that can both produce and consume energy, e.g., residential users with solar panels or electric vehicles. Prosumers have both *unresponsive* and *responsive* loads. Responsive loads, such as heating, ventilation, and air conditioning (HVAC) systems can adjust their load to reduce costs (e.g., store energy in thermal form anticipating high energy prices). On the contrary, unresponsive loads do not change their consumption regardless of the prices (the flexibility of loads can change throughout the day).

Prosumers express their intention (and conditions) to trade energy through bids. We represent a bid as the following tuple:

$$\langle \tau, \sigma, \pi \rangle,$$

where τ specifies the time interval in which energy exchange can occur; σ indicates the maximum amount of energy available to trade; and π denotes the reservation price (minimum or maximum price accepted by sellers or buyers, respectively).

We assume that prosumers cannot change their bills by tampering the meters that measure their physical energy flow.

2) *Blockchain based Electricity Market*: A blockchain is a distributed ledger, this means that multiple nodes have a copy of the transactions. Special nodes (called miners) decide the state of the distributed ledger (e.g., the transactions) through a consensus protocol, which induces a high cost to modify the ledger e.g., *Proof of Work* (PoW) [7], [25] and *Proof of Stake* (PoS) [26]. The blockchain creates a chain-like data structure in which each block has a reference to previous blocks; in this way, the transactions recorded become practically immutable. Thus, blockchains provide trustworthy data storage and computation (in the form of smart contracts) without requiring a trusted entity.

3) *Gateways*: Prosumers may not participate directly in the blockchain network, because consensus protocols typically have high computational and storage requirements, which IoT based energy trading devices cannot satisfy. Hence, prosumers may access the distributed energy market through *gateway* nodes. A gateway either forwards messages between the prosumers and the distributed energy market or acts as miner, which execute the blockchain consensus protocol. To protect the prosumers' privacy, the communication between prosumers and gateways may be encrypted and anonymized, as described in [27]. Gateways can be operated by the company that implements the TES or by a third party.

4) *Distribution System Operator (DSO)*: Besides the information infrastructure, the system needs a continuous management of the physical infrastructure. In this case, we assume that a DSO supervises the system and is responsible for managing the distribution grid, billing, installing smart meters, satisfying the net demand, and maintaining stability [28]. Although we refer to the DSO as the system's manager, other entities, such as electric utilities, can be in a better position to provide these services.

B. Electricity Market

Power systems use electricity markets to find efficient and reliable operations. Efficiency involves maximizing the benefit of all the participants, while reliability refers to satisfying engineering constraints. Electricity markets operate periodically (gather bids and determine the system's operation) to guarantee a correct operation at any moment. In our notation, we omit the time when the transactions occur; however, we reiterate that the market's operation occurs periodically (*e.g.*, the market accepts bids to decide the price and trades every five minutes).

1) *Ideal Model*: Let us denote the set of prosumers as \mathcal{P} . We classify prosumers in two sets, namely the set of consumers \mathcal{C} and the set of generators \mathcal{G} . We define the utility of a consumer $i \in \mathcal{C}$ as

$$u_i(q_i, p) = v_i(q_i) - q_i \cdot p,$$

where $v_i(q_i)$ represents the benefit obtained from consuming $q_i \geq 0$ units of energy, and $p \in \mathbb{R}$ is the unit price paid for the energy. We define the profit of a generator $j \in \mathcal{G}$ as its income minus its generation cost, that is,

$$u_j(q_j, p) = q_j \cdot p - C_j(q_j),$$

where $C_j(q_j)$ represents the cost of producing $q_j \geq 0$ units of energy. We assume that the market decides the trades within a feeder. Therefore, consumption q_i represents energy consumed by the feeder's loads. Moreover, production q_j represents energy produced within the feeder or energy transferred by external sources to supply loads within the feeder. Thus, local generators can export energy outside the feeder, but the local market is not involved in these trades.

Ideally, the power system distributes resources efficiently, that is, maximizing the benefit of all the participants. In this

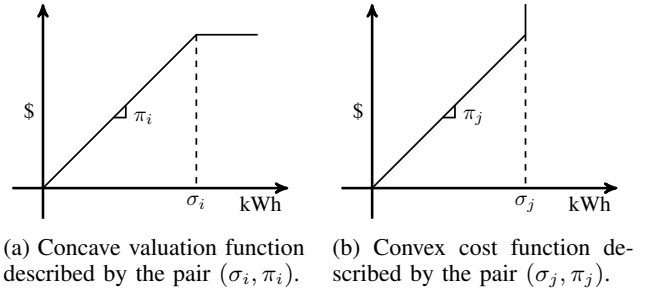


Fig. 2: Example of piecewise linear functions.

case, we use the *social welfare* as our efficiency metric, which is defined as

$$f(\mathbf{q}, p) = \sum_{i \in \mathcal{C}} u_i(q_i, p) + \sum_{j \in \mathcal{G}} u_j(q_j, p),$$

where the vector $\mathbf{q} = [q_i]_{i \in \mathcal{P}}$ denotes the energy consumed and generated by the prosumers.

Power markets must maintain a balance between trades; in other words, the total energy sold must equal the total energy bought. In practice, a small part of the energy delivered may dissipate as heat in transmission lines. If we ignore these transmission losses, we can express the balance condition as

$$\sum_{i \in \mathcal{C}} q_i = \sum_{j \in \mathcal{G}} q_j.$$

A system that operates efficiently and reliably has to allocate resources maximizing social welfare and satisfying the system's physical constraints. We formulate the system's goal as the solution of the following optimization problem:

$$\underset{\mathbf{q}, p}{\text{maximize}} \quad f(\mathbf{q}, p) \quad (1a)$$

$$\text{subject to} \quad \sum_{i \in \mathcal{C}} q_i = \sum_{j \in \mathcal{G}} q_j, \quad (1b)$$

$$p = \partial C_j / \partial q_j(q_j), \quad \forall j \in \mathcal{G} \quad (1c)$$

Eq. (1b) captures the need to balance the energy transactions, while Eq. (1c) captures the optimal production of generators (marginal costs equal the energy price). Let us denote the optimal equilibria as (\mathbf{q}^*, p^*) , where $\mathbf{q}^* = [q_i^*]_{i \in \mathcal{P}}$ represents the energy traded (production/consumption) and p^* denotes the market's clearing price.

Although the market defines a unique price, consumers also pay for congestion and transmission losses. For this reason, the price paid by a prosumer, which is also called locational marginal price (LMP), changes with its location. Systems without transmission constraints and losses have a single LMP; hence, prosumers observe the same price in our model.

Electricity markets often solve Eq. (1a) through auctions, in which the participants have incentives to reveal their private information *i.e.*, the functions $v_i(\cdot)$ and $C_j(\cdot)$ [29]. Moreover, economic models typically assume that the valuation function $v_i(\cdot)$ is concave and that the cost function $C_j(\cdot)$ is convex. As a consequence, Eq. (1a) has a unique solution.

2) *Approximated Model*: Market operators often restrict the form of the bids [30]. For example, some auctioneers assume that the functions $v_i(\cdot)$ and $C_j(\cdot)$ belong to some family of functions. In this work, we assume that each prosumer i has a piece-wise linear function defined by a pair (σ_i, π_i) , where σ_i denotes the maximum energy available (or needed) and π_i specifies the minimum (or maximum) unit price accepted by the prosumer [31], [32]. Fig. 2 shows an example of the functions supported in our market.

With this simplified model we can use a *double-auction* to solve the optimization problem in Eq. (1a). The double auction creates demand and offer curves to find the market’s equilibria. The demand curve is a piecewise linear function constructed ordering the buyers’ bids in descending order by their prices. Likewise, the offer curve is constructed ordering the sellers’ bids in ascending order by their prices. The intersection of these curves yields the market’s clearing price and the total energy traded in the market.

Since the market must account for the total demand, it has to estimate the unresponsive loads. For simplicity, we estimate the unresponsive loads as the difference between the total amount of energy demanded and the total load. In our model, we assume that the DSO makes this estimate and submits a bid requesting that amount of energy at the maximum price allowed [31]. Note that other entities could also act in this role; our analysis and experimental results are not contingent on this assumption.

Using a double auction ensures that there is no consumption or production strategy that will result in better utility for the prosumers than truthfully reporting σ_i, π_i , which characterize the functions $v_i(\cdot)$ and $C_j(\cdot)$. However, in Section V, we show that adversarial prosumers may profit by affecting other prosumers’ bids.

3) *Distributed Market Operation*: The operation of the distributed market is ensured by the blockchain miner nodes, and it is often implemented in the form of a smart contract (we refer the reader to [22] for more details). Typically, the operation of the market follows the next iterative steps:

- 1) The market starts in the *Bidding* state for a future time interval τ .
- 2) Prosumers submit their bids for the interval τ to the gateways, which register the bids on the market (blockchain).
- 3) Market transits to *Clearing* state for interval τ . In this case, we implement a double auction [31], [33], which finds transactions that maximize social welfare. Nonetheless, smart contracts provide enough flexibility to implement other types of auctions and restrictions. The gateways forward the transactions and clearing price to the prosumers, who will later exchange energy accordingly (recall Section III-B2 and Fig. 2).
- 4) Market transits to *Bidding* state for time $\tau + 1$, and the gateways notify the prosumers of the beginning of the next trading interval.

TABLE I: Adversary’s Attack Scenarios

Scenario	Know bids prior to the attack	Target individual bidders
1	✓	✓
2	x	✓
3	x	x

IV. THREAT MODEL

In this section we describe the adversary’s capabilities, its goal, and attack strategies.

A. Adversary’s Capabilities

We assume that the adversary cannot tamper with or remove bids accepted by the market, and it cannot tamper with or disrupt the market clearing mechanism (the blockchain guarantees that this requires a large amount of resources). However, the adversary—who may be one of the prosumers—can read past bids and clearing prices from the blockchain.

Blockchains can suffer from several vulnerabilities, some of which lead to thefts of cryptocurrencies or public keys [34]. An adversary may leverage these vulnerabilities to tamper with the prosumers’ bids. For example, an adversary may steal the public keys of prosumers to forge bids or compromise smart appliances or transactive controllers to modify their bidding strategies. However, it may be much easier to compromise a single node that is acting as a gateway for a group of prosumers, than attacking multiple prosumers individually. For example, the adversary can exploit bugs in the Ethereum software to either bypass authentications or to disable miners [35]. We consider three attack scenarios against miners that differ in the adversary’s knowledge and capabilities of (see Table I for a summary).

- 1) *Gateway Confidentiality and Integrity Attack*: The adversary compromises a gateway and obtains sufficient access to delay or discard particular bids (*i.e.*, prevent them from being recorded on the market). In this scenario, the adversary is also capable of reading all bids before deciding which bids to discard (*e.g.*, by reading the bids submitted to the compromised gateway as well as the ones recorded on the blockchain by other gateways).
- 2) *Gateway Integrity Attack*: The adversary can discard or delay selected bids; however, the adversary must decide which bids to discard without complete information, relying only on historical data about the prosumers’ past bids.
- 3) *Gateway Availability Attack*: The adversary cannot delay particular bids, but it has sufficient resources to launch a DDoS attack against one of the gateways. This attack prevents the processing of some bids in the market, but the adversary cannot read bids either.

B. Adversary’s Goal

We consider a rational, profit-oriented adversary, who is interested in maximizing its own profit. The adversary’s goal and strategy depend on its role (*e.g.*, generator or consumer). We focus our attention on *adverse generators*, who discard

(or delay) the bids of prosumers. Concretely, we assume that adverse generators pursue a market equilibria (q^a, p^a) that increases the generator's profit by $\lambda\%$. We express this condition as

$$\sum_{j \in \mathcal{G}} u_j(q_j^a, p^a) = \sum_{j \in \mathcal{G}} (1 + \lambda) u_j(q_j^*, p^*). \quad (2)$$

where \mathcal{G} represents the set of adverse generators.

Let the total generation before and after the attack be

$$Q^* = \sum_{j \in \mathcal{G}} q_j^* \quad \text{and} \quad Q^a = \sum_{j \in \mathcal{G}} q_j^a,$$

respectively. Now, let us approximate the aggregate cost function with the following quadratic function (this is a common approximation in the literature [36]):

$$\sum_{j \in \mathcal{G}} C_j(q_j) \approx C(Q) = \beta_2 \cdot Q^2 + \beta_1 \cdot Q,$$

where $Q \geq 0$ and β_1 and β_2 are constants. Thus, the aggregate utility of sellers without an attack is

$$\sum_{j \in \mathcal{G}} u_j(q_j^*, p^*) = Q^* \cdot p^* - C(Q^*). \quad (3)$$

Since $p^* = \dot{C}(Q^*)$, we can rewrite Eq. (3) as

$$\sum_{j \in \mathcal{G}} u_j(q_j^*, p^*) = \beta_2 (Q^*)^2. \quad (4)$$

Now, substituting Eq. (4) into Eq. (2), we get

$$Q^a = \sqrt{1 + \lambda} Q^*.$$

Thus, the adversary attempts to increase the total energy traded in the market by

$$\Delta Q^a = Q^* (\sqrt{1 + \lambda} - 1). \quad (5)$$

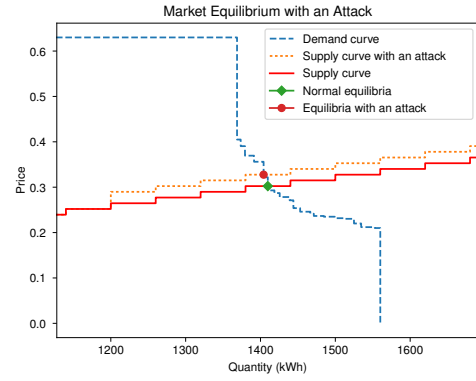
We assume that the adversary chooses λ subject to some constraints. First, the adversary limits the impact of its attacks to avoid damaging the system, which may affect its own assets and future gains. Second, the market behavior is monitored by regulators for sudden or excessive deviation from the expected values. When detected, these deviations may lead to investigations, which can result in the discovery of the attack and the punishment of the adversary. Regulators already monitor the performance of firms to punish those that exercise market power [37].

V. ANALYSIS

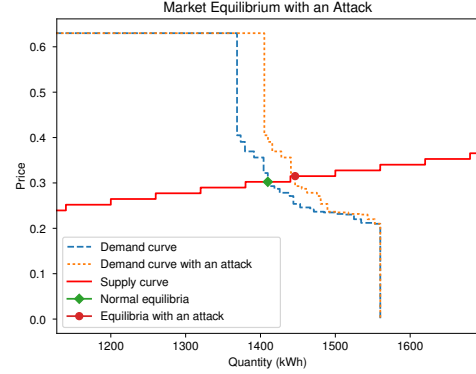
In this section, we discuss strategies that an adversary may use to increase its profit, given the capabilities that we assumed in Section IV-A. Then, we discuss strategies for mitigating such attacks.

A. Attack Strategy

An adverse generator benefits from delaying the bids of prosumers if the resulting market equilibrium increases its profit. Fig. 3 shows how a delay attack on either buying or selling bids changes the equilibrium. In a double auction, the offer and demand curves capture the trades (price and quantity) that buyers and sellers would accept. Their intersection corresponds to the market equilibrium, a condition



(a) Market equilibria when delaying the bids of generators.



(b) Market equilibria when delaying the bids of consumers.

Fig. 3: An adverse generator can increase the market's equilibria price delaying bids of both buyers and sellers.

in which no prosumer would change its trades. Delays in bids of competing generators (who offer lower prices) forces the market to procure energy from more expensive generators, which raises the prices (Fig. 3a illustrates this). We leave the analysis of such attacks to future work.

The demand curve is constructed with bids ordered by descending price. In our case, the DSO constructs bids for the estimated unresponsive loads, which accept the maximum price allowed in the market. The demand curves in Fig. 3 have flat regions corresponding to bids of unresponsive loads. The decreasing regions correspond to the bids of responsive loads.

Delays in buyers' bids can also benefit the adversary, because missing bids may lead to overestimation of the unresponsive loads. In other words, the DSO may assume that the appliances that do not submit bids will accept any price. In such cases, the demand curve changes reflecting a higher willingness to pay for energy, which raises the prices (see Fig. 3b). Next, we analyze this attack in the three scenarios that we introduced in Section IV-A.

1) *Confidentiality and Integrity Attack*: In this scenario, the adversary can collect all the bids submitted to the compromised gateway, and read the bids submitted to the other gateways. Hence, it can compute the market's clearing price p^* and the total energy traded Q^* . The adversary uses these

values to calculate the desired deviation in the trades ΔQ^a (see Eq. (5)). Then, it selects a subset of bids \mathcal{V} such that

$$\sum_{v \in \mathcal{V}} q_v^* \approx \Delta Q^a.$$

In practice the impact of the attack will be lower than ΔQ^a , because some appliances may reduce their load as a response to higher prices.

This is an ideal scenario for the adversary, since it is able to discard bids possessing complete information about the market. Further, this scenario allows the adversary to determine which gateway is the optimal target for the attack.

2) *Integrity Attack*: In this scenario, the adversary cannot observe the bids before they reach the market. Instead, the adversary must rely on historical information about the bids in order to choose its attack strategy, since it can observe only past bids and clearing prices. Hence, the adversary must select which gateway to attack and which consumers' bids to discard based on historical data and trends.

Since the total demand typically follows daily patterns, we assume that it is possible to predict whether the bid of a prosumer will fall below the market price. Let us model the bids and the clearing price for a particular time interval as random variables. We denote with the random variable p^τ the market clearing price, and let $(\sigma_i^\tau, \pi_i^\tau)$ denote random variables of prosumer i 's bid in the time interval τ .

We use the random variable

$$x_i^\tau = \begin{cases} 1 & \text{if } \pi_i^\tau \leq p^\tau \\ 0 & \text{otherwise.} \end{cases}$$

to denote whether discarding the bid of consumer i changes the market equilibria.

Let us denote with ϱ_i^τ the probability that the adversary delays the bid of consumer i during the time interval τ . Further, we denote the adversary's strategy during timer interval τ as $\varrho^\tau = [\varrho_i^\tau]_{i \in \mathcal{C}}$. We approximate the impact of an attack (the reduction in demand) as

$$\Gamma(\varrho^\tau) = \sum_{i \in \mathcal{C}^a} \varrho_i^\tau \mathbb{E}[x_i^\tau \cdot q_i^\tau],$$

where \mathcal{C}^a denotes the set of consumers that use the gateway targeted in the attack. In this case, the adversary chooses its attack strategy ϱ^τ to achieve the desired impact during time interval τ , which we denote as $\Delta Q^{a,\tau}$; that is,

$$\Gamma(\varrho^\tau) = \Delta Q^{a,\tau}. \quad (6)$$

In particular, by selecting the bids with the highest expected impact the adversary minimizes the number of targets. Without loss of generality let us rank the bids according to their expected impact. Concretely, let $\mathbb{E}[x_i^\tau \cdot q_i^\tau] \geq \mathbb{E}[x_j^\tau \cdot q_j^\tau]$ if $i > j$, for all bids i and j . Now, the adversary can select with probability $\varrho_i^\tau = 1$ the first m bids that satisfy

$$\sum_{i=1}^m \mathbb{E}[x_i^\tau \cdot q_i^\tau] \leq \Delta Q^{a,\tau} \leq \sum_{i=1}^{m+1} \mathbb{E}[x_i^\tau \cdot q_i^\tau]$$

and select the $m + 1$ bid with a probability ϱ_{m+1}^τ that satisfies Eq. (6).

3) *Availability Attack*: In this scenario, the adversary cannot delay particular bids. Instead, the adversary launches a DoS attack which leads to discarding bids randomly, with uniform probability. Hence, $\varrho_j^\tau = \varrho^\tau = \varrho_i^\tau$ for all prosumers i and j that use the targeted gateway, where ϱ^τ depends on the intensity of the attack. Similar to the previous scenario, the adversary must rely on historical data to select the gateway and choose the intensity of the attack. Thus, adversary selects

$$\varrho^\tau = \Delta Q^{a,\tau} / \sum_{i \in \mathcal{C}^a} \mathbb{E}[x_i^\tau q_i^\tau].$$

We assume that this attack does not affect the bids of the adversary, who can use a different gateway if necessary.

B. Mitigation Strategy

If a consumer does not receive confirmation from the gateway, then it might submit its bid to another gateway, after waiting a certain amount of time.

The amount of time that is required to record a transaction (e.g., a bid) on a blockchain depends on the time required to generate the next block, which is non-deterministic, and for most blockchains follows an exponential distribution [38]. Thus, consumers can wait until they have some confidence about the state of the gateway (e.g., waiting for two standard deviations of the block generation time) before resubmitting.

Frequent resubmissions may harm the performance of gateways; hence, it is possible that a resubmitted bid is not recorded in time for inclusion in the market clearing. For this reason, prosumers may select wait times according to their risk preferences. Alternatively, prosumers may reduce the efficacy of attacks by selecting gateways randomly, thereby increasing the adversary's uncertainty about which gateway to target.

VI. EXPERIMENTAL EVALUATION

A. Testbed Implementation

For experimental evaluation, we deployed GridLAB-D [39] and a private Ethereum blockchain network [2]. GridLAB-D simulates the smart grid, including prosumer logic for creating bids. GridLAB-D models retail markets through double auctions [31] that run every five minutes. Our power system has 58 residential commercial houses, which in turn incorporate appliances such as heating, ventilation, and air conditioning (HVAC) systems. GridLAB-D models the response of the loads to weather and market's prices, giving realism to the simulations. In this case, *transactive controllers* manage HVAC systems and make bids in the market.

The blockchain stores bids, market clearing prices, and calculates the market equilibria with a smart contract. We built our testbed on the open-source TRANSAX framework [22], which provides the prosumer interfaces and a smart contract. Each prosumer is assigned to one of three Ethereum clients, which act as gateways to the private Ethereum network. Based on the prosumers' allocation, the attacker chooses one of the Ethereum clients to attack, and delays a subset of the bids sent to that client. Since each bid is valid for a single interval, this in effect discards the bids.

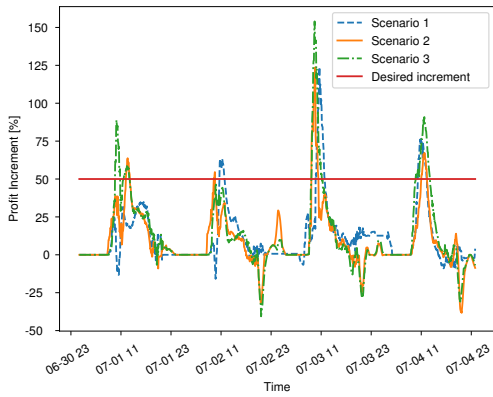


Fig. 4: Impact of attacks (snapshot of five days, from July 1 at 00:00 to July 5 at 00:00).

B. Experimental Results

In the simulations, we use bids and the market's equilibria without attacks (q^*, p^*) and with them (q^a, p^a) to calculate the surplus of sellers (see Eq. (3)); for instance, we calculate the profit without attacks as

$$\sum_{j \in \mathcal{G}} u_j(q_j^*, p^*) = Q^* p^* - \sum_{j \in \mathcal{G}} \pi_j q_j^*.$$

We measure the impact of the attack as the profit increment for the generators:

$$\text{Profit increment} = \frac{\sum_{j \in \mathcal{G}} \{u_j(q_j^a, p^a) - u_j(q_j^*, p^*)\}}{\sum_{j \in \mathcal{G}} u_j(q_j^*, p^*)}$$

Fig. 4 shows the gains of sellers with attack scenarios described in Table I. This figure illustrates that the attacks have a larger impact between 9am and 1pm, when some of the responsive loads (e.g., air conditioning systems) operate. Likewise, periods with a low or even negative benefit coincide with periods with less market activity. High prices caused by the attack alter the demand patterns of some smart appliances. For this reason, the attack's actual impact exceeded the desired gains (we designed the attacks to increase the profit by at most 50%). We also see that for Scenarios 2 and 3, producers can experience losses, which typically occur around 8pm.

Next, we investigate the impact of attack as a function of the mitigation rate. In this case, the mitigation rate denotes the ability of prosumers to submit their bids despite an attack on their gateway. For example, prosumers who use only the targeted gateway have a mitigation rate of 0%.

Figs. 5 to 7 show distributions of the attack impact for the three attack scenarios with different mitigation rates. These figures show that higher mitigation rates tend to reduce the impact of attacks (both their mean and their standard deviation). The impact is highly variable, but on average, they improve the generators' utility. Scenario 3, which requires the least privileges, suffers from higher variance but still provides increased utility. Scenario 2 has lower variance, but also achieves lower impact than the other scenarios.

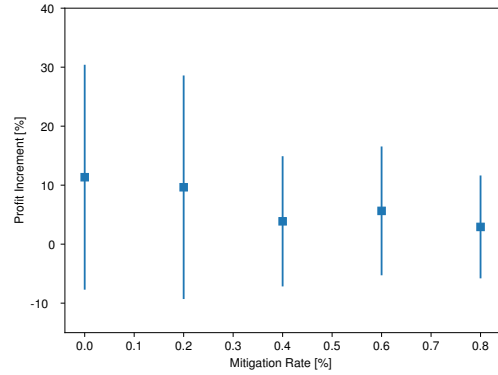


Fig. 5: *Gateway Confidentiality and Integrity Attack*: The adversary designs its attack knowing the bids.

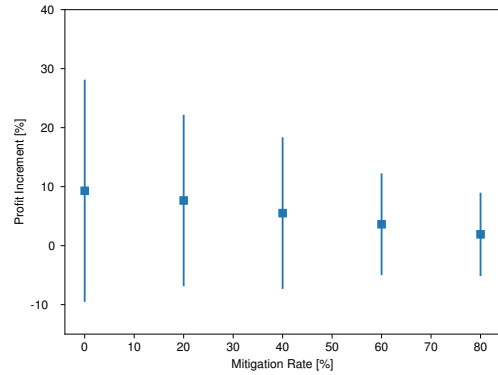


Fig. 6: *Gateway Integrity Attack*: The attacker chooses a gateway to attack based on historical prosumer data. Then, bids from the targeted prosumers are discarded as they arrive.

Our private blockchain network takes 11 s to generate blocks (on average), with a standard deviation of 11.4 s. We assume that the block generation time follows an exponential distribution; hence, we can compute the likelihood of a false-positive alert (i.e., timeout without an attack) as $F(x; \eta) = 1 - e^{-\eta x}$. Assuming that a prosumer waits for two standard deviations, the wait time is 33 s. Since $1 - e^{-\frac{1}{11} \times 33} = 0.95$, the prosumer will have a 5% false-positive rate. Since bids have to be recorded on the blockchain in the first 4 minutes of the interval, many blocks will likely be mined before clearing. In fact, a prosumer may attempt to resubmit around 7 times.

VII. CONCLUSION

In this paper, we examined blockchain based TES, which have received significant attention recently due to their unique advantage of providing resilience and integrity in decentralized systems. We introduced a novel class of cyber-attacks, which do not target the trading system directly, but rather target the interface between the prosumers and the system. We found that even simpler attacks, such as (D)DoS, can effectively manipulate the clearing price of a blockchain based market. However, we have also demonstrated that the threat can be mitigated via detection and gateway switching. We evaluated

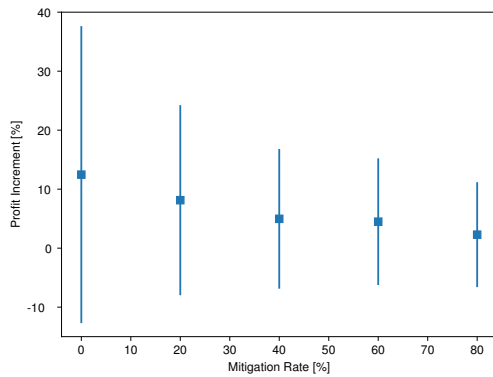


Fig. 7: *Gateway Availability Attack*: Attacker chooses a gateway and executes a (D)DoS attack, resulting in some random subset of bids being discarded.

the impact of these attacks experimentally using a testbed based on GridLAB-D and a private Ethereum network.

In the future, we will extend our analysis to consider proactive defenses (e.g., through random selection of gateways), more sophisticated attack detection, and cyber-attacks on individual prosumers (i.e., compromising their IoT devices).

ACKNOWLEDGMENT

This work has been supported in part by the National Science Foundation under award numbers 1818901 and 1840052 and by the National Institute of Standards and Technology under Grant 70NANB18H198. We also thank the anonymous reviewers of our submission.

REFERENCES

- [1] K. Kok and S. Widergren, "A society of devices: Integrating intelligent distributed resources with transactive energy," *IEEE Power and Energy Magazine*, vol. 14, no. 3, pp. 34–45, 2016.
- [2] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," Ethereum Project – Yellow Paper, Tech. Rep. EIP-150, April 2014.
- [3] A. Mavridou, A. Laszka, E. Stachtari, and A. Dubey, "VeriSolid: Correct-by-design smart contracts for Ethereum," in *23rd International Conference on Financial Cryptography and Data Security (FC)*, February 2019, pp. 446–465.
- [4] N. Wang, X. Zhou, X. Lu, Z. Guan, L. Wu, X. Du, and M. Guizani, "When energy trading meets blockchain in electrical power system: The state of the art," *Applied Sciences*, vol. 9, no. 8, p. 1561, 2019.
- [5] M. Mylrea and S. N. G. Gourisetti, "Blockchain for smart grid resilience: Exchanging distributed energy at speed, scale and security," in *2017 Resilience Week (RWS)*. IEEE, 2017, pp. 18–23.
- [6] D. P. Chassin, K. Schneider, and C. Gerkensmeyer, "GridLAB-D: An open-source power systems modeling and simulation environment," in *2008 IEEE/PES Transmission and Distribution Conference and Exposition*. IEEE, 2008, pp. 1–5.
- [7] V. Buterin *et al.*, "Ethereum white paper," 2013.
- [8] K. Zetter, "Inside the cunning, unprecedented hack of Ukraine's power grid," *WIRED Magazine*, March 2016. [Online]. Available: <https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/>
- [9] H. He and J. Yan, "Cyber-physical attacks and defences in the smart grid: a survey," *IET Cyber-Physical Systems: Theory & Applications*, vol. 1, no. 1, pp. 13–27, 2016.
- [10] L. Jia, R. J. Thomas, and L. Tong, "Malicious data attack on real-time electricity market," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2011, pp. 5952–5955.
- [11] L. Xie, Y. Mo, and B. Sinopoli, "Integrity data attacks in power market operations," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 659–666, 2011.
- [12] Y. Chen, Y. Tan, and B. Zhang, "Exploiting vulnerabilities of load forecasting through adversarial attacks," in *10th ACM International Conference on Future Energy Systems (e-Energy)*, 2019, pp. 1–11.
- [13] C. Barreto and X. Koutsoukos, "Design of load forecast systems resilient against cyber-attacks," in *10th Conference on Decision and Game Theory for Security (GameSec)*, 2019, pp. 1–20.
- [14] C. Barreto and A. Cardenas, "Impact of the market infrastructure on the security of smart grids," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2018.
- [15] C. Barreto and X. Koutsoukos, "Attacks on electricity markets," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019.
- [16] B. Krebs, "Who is Anna-Senpai, the Mirai Worm Author?" Krebs on Security, 2017, accessed: October 23rd, 2019. [Online]. Available: <https://krebsonsecurity.com/2017/01/who-is-anna-senpai-the-mirai-worm-author/>
- [17] I. Stelliou, P. Kotzanikolaou, M. Psarakis, C. Alcaraz, and J. Lopez, "A survey of IoT-enabled cyberattacks: Assessing attack paths to critical infrastructures and services," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 3453–3495, 2018.
- [18] W. Chin, W. Li, and H. Chen, "Energy big data security threats in iot-based smart grid communications," *IEEE Communications Magazine*, vol. 55, no. 10, pp. 70–75, Oct 2017.
- [19] S. Soltan, P. Mittal, and H. V. Poor, "BlackIoT: IoT botnet of high wattage devices can disrupt the power grid," in *27th USENIX Security Symposium*. Baltimore, MD: USENIX Association, 2018, pp. 15–32.
- [20] L. Energy, "Exergy-building a robust value mechanism to facilitate transactive energy," *LO3. Retrieved March*, vol. 29, p. 2019, 2017.
- [21] "Power Ledger whitepaper," PowerLedger, July 2017. [Online]. Available: <https://powerledger.io/whitepaper/>
- [22] A. Laszka, S. Eisele, A. Dubey, G. Karsai, and K. Kvaternik, "TRANSAX: A blockchain-based decentralized forward-trading energy exchange for transactive microgrids," in *Proceedings of the 24th IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, December 2018, pp. 918–927.
- [23] A. Laszka, A. Dubey, M. Walker, and D. Schmidt, "Providing privacy, safety, and security in IoT-based transactive energy systems using distributed ledgers," in *7th International Conference on the Internet of Things (IoT)*, October 2017, pp. 13:1–13:8.
- [24] A. Wörner, A. Meeuw, L. Ableitner, F. Wortmann, S. Schopfer, and V. Tiefenbeck, "Trading solar energy within the neighborhood: field implementation of a blockchain-based electricity market," in *Proceedings of the 8th DACH+ Conference on Energy Informatics*, vol. 2, 2019.
- [25] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," Working Paper, 2008.
- [26] S. King and S. Nadal, "Ppcoin: Peer-to-peer crypto-currency with proof-of-stake," *self-published paper*, August, vol. 19, 2012.
- [27] J. Bergquist, A. Laszka, M. Sturm, and A. Dubey, "On the design of communication and transaction anonymity in blockchain-based transactive microgrids," in *1st Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers (SERIAL)*. ACM, 2017, pp. 3:1–3:6.
- [28] O. Dag and B. Mirafzal, "On stability of islanded low-inertia microgrids," in *2016 Clemson University Power Systems Conference (PSC)*. IEEE, 2016, pp. 1–7.
- [29] R. B. Myerson, "Optimal auction design," *Mathematics of operations research*, vol. 6, no. 1, pp. 58–73, 1981.
- [30] R. Baldick, "Electricity market equilibrium models: The effect of parametrization," *IEEE Transactions on Power Systems*, vol. 17, no. 4, pp. 1170–1176, 2002.
- [31] J. C. Fuller, K. P. Schneider, and D. Chassin, "Analysis of residential demand response and double-auction markets," in *2011 IEEE Power and Energy Society General Meeting*, July 2011, pp. 1–7.
- [32] S. Li, W. Zhang, J. Lian, and K. Kalsi, "Market-based coordination of thermostatically controlled loads—Part I: A mechanism design formulation," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1170–1178, 2016.
- [33] E. Mengelkamp, J. Gärtner, and C. Weinhardt, "Decentralizing energy systems through local energy markets: the LAMP-project," in *Multikonferenz Wirtschaftsinformatik*, 2018.

- [34] H. Chen, M. Pendleton, L. Njilla, and S. Xu, "A survey on Ethereum systems security: Vulnerabilities, attacks and defenses," *arXiv preprint arXiv:1908.04507*, 2019.
- [35] C. Details, "Ethereum: Vulnerability statistics," CVE Details, 2020. [Online]. Available: <https://www.cvedetails.com/vendor/17524/Ethereum.html>
- [36] D. S. Kirschen and G. Strbac, *Fundamentals of power system economics*. John Wiley & Sons, 2018.
- [37] J. Lazar, F. Weston, W. Shirley, J. Migden-Ostrander, D. Lamont, and E. Watson, "Revenue regulation and decoupling: A guide to theory and application," Regulatory Assistance Project, Tech. Rep., 2016. [Online]. Available: <https://www.raonline.org/knowledge-center/revenue-regulation-and-decoupling-a-guide-to-theory-and-application-incl-case-studies/>
- [38] R. Bowden, H. P. Keeler, A. E. Krzesinski, and P. G. Taylor, "Block arrivals in the bitcoin blockchain," *arXiv*, vol. abs/1801.07447, 2018.
- [39] J. Fuller, "Transactive modeling and simulation capabilities," in *NIST Transactive Energy Challenge Preparatory Workshop*, 2015.