

Safe and Private Forward-Trading Platform for Transactive Microgrids

SCOTT EISELE, Vanderbilt University
TAHA EGHTEHAD, University of Houston
KEEGAN CAMPANELLI, Vanderbilt University
PRAKHAR AGRAWAL, Vanderbilt University
ARON LASZKA, University of Houston
ABHISHEK DUBEY, Vanderbilt University

Transactive microgrids have emerged as a transformative solution for the problems faced by distribution system operators due to an increase in the use of distributed energy resources and rapid growth in renewable energy generation. Transactive microgrids are tightly coupled cyber and physical systems, which require resilient and robust financial markets where transactions can be submitted and cleared, while ensuring that erroneous or malicious transactions cannot destabilize the grid. In this paper, we introduce TRANSAX, a novel decentralized platform for transactive microgrids. TRANSAX enables participants to trade in an energy futures market, which improves efficiency by finding feasible matches for energy trades, reducing the load on the distribution system operator. TRANSAX provides privacy to participants by anonymizing their trading activity using a distributed mixing service, while also enforcing constraints that limit trading activity based on safety requirements, such as keeping power flow below line capacity. We show that TRANSAX can satisfy the seemingly conflicting requirements of efficiency, safety, and privacy, and we demonstrate its performance using simulation results.

CCS Concepts: • **Computer systems organization** → *Embedded and cyber-physical systems*; • **Security and privacy**;

Additional Key Words and Phrases: smart grid, transactive energy, distributed ledger, privacy, decentralized application, cyber-physical system, smart contract, blockchain

1 INTRODUCTION

The traditional setup of the power grid is rapidly changing. Solar panel capacity is estimated to grow from 4% in 2015 to 29% in 2040 [42], and with the decreasing costs of battery technology, it is becoming increasingly feasible to support almost 99% of total load with renewable sources, by balancing out the intermittence with batteries [29]. These changes are also leading to the development of a decentralized vision for the future of power-grid operations, in which local peer-to-peer energy trading within microgrids can be used to reduce the load on distribution system operators (DSO), leading to the development of transactive energy systems (TES) [10, 25, 35, 41]. A transactive energy system is a set of market based constructs for dynamically balancing the demand and supply across the electrical infrastructure [35]. In this approach, customers connected by transmission and distributions lines can participate in an open market, trading and exchanging energy locally. Customers participating in these markets are known as *prosumers*. There are typically three phases in the operation of this market: posting offers to buy or sell energy, matching selling offers with buying offers, and synchronized energy transfer to and from the grid.

In theory, these interactions could happen in a centralized manner by communicating all offers to a centralized market, which would match the offers and broadcast the trades back to the individual

Authors' addresses: Scott Eisele, scott.r.eisele@vanderbilt.edu, Vanderbilt University, Nashville, TN; Taha Eghtesad, teghtesad@uh.edu, University of Houston, Houston, TX; Keegan Campanelli, keegan.m.campanelli@vanderbilt.edu, Vanderbilt University; Prakhar Agrawal, , Vanderbilt University; Aron Laszka, alaszka@houston.edu, University of Houston; Abhishek Dubey, Vanderbilt University, abhishek.dubey@vanderbilt.edu.

prosumers. However, in a centralized solution [23], the market presents a single point of failure. Building a decentralized market for transactive microgrids is challenging because the system has to satisfy a number of requirements. First, the market must be **efficient**: the system should maximize the utilization of local supply—in meeting local demand—by matching the prosumers in the microgrid, taking advantage of their temporal flexibility in production and consumption. This requirement is crucial since effective trading is the purpose of the system; all other requirements are supporting this one. Second, the market must be **safe**: the system must prevent negligent or malicious trading from endangering the stability or physical safety of the microgrid by rejecting trades that are unlikely to be delivered or would violate line capacity constraints.¹ Third, the market must be **private**: the system should conceal information that could be used to infer the prosumers’ energy usage patterns. The amount of energy produced, consumed, bought, or sold by any prosumer should be anonymous to other prosumers and limited to the monthly bill for the DSO. This is necessary because such information could be exploited, e.g., to determine when a resident is at home. Fourth, the market must be **secure**: the system must ensure authenticity, data integrity, and auditability for offers, trades, and bills. Finally, the market must be **resilient**: it must retain availability even if some nodes or entities (e.g., DSO) are unavailable.

The research community is increasingly advocating the use of distributed ledgers in the energy sector [1]. This is primarily because a distributed ledger can provide an immutable, complete, and fully auditable record of all transactions that have occurred within a system. However, there are still research gaps. For example, Adoni et al. [1] surveyed 140 applications of distributed ledgers and found that 33% were focused on decentralized energy trading, with privacy preservation being one of the key challenges that has not been addressed. They also highlighted balancing supply to demand (stability) as another critical issue. Some works validate transactions using the DSO to ensure that the decentralized energy trades agreed between parties can actually take place. However, doing so comes at a loss of privacy because DSO is made aware of each trade. Brenzikofer et al. consider privacy mechanisms [7]. However, they only incentivize stability through dynamic grid tariffs, and do not actively enforce stability in the sense of limiting trading that can destabilize the grid. The work presented in this paper addresses the problem of privacy while ensuring the system’s safety can be enforced.

Contributions: In this paper, we introduce TRANSAX², a distributed-ledger based transactive energy system. In contrast to prior work, the key innovation of TRANSAX is addressing the requirements of safety, efficiency, privacy, security, and resilience together. The specific contributions of this paper are as follows.

To ensure trading safety, we introduce energy production and consumption assets that represent trading allowances. Based on these assets, we enforce constraints, which can correspond to line capacity constraints, on individual prosumers as well as on groups of prosumers (e.g., prosumers that are connected to the same feeder). Note that the trading safety requirement is not difficult to satisfy by itself; it poses a challenge when combined with the privacy requirement. To provide privacy, we integrate a mixer that enables prosumers to anonymize their production and consumption assets within their groups, and hence trade anonymously. We show that prosumers can remain anonymous within their groups, while the system can still enforce safety constraints (e.g., feeder line capacity) by transforming them into group constraints. We also introduce privacy-preserving billing that enables prosumers to disclose their trades only to their smart meters, which can then perform billing without disclosing anything to the DSO other than the billed amount.

¹ Note that this is orthogonal to the *physical enforcement of safety* which is provided by overcurrent protection units that limit the total current flowing through the microgrid. ² This paper is a significant extension of our previously published conference papers [14, 26, 27].

To improve trading efficiency, we provide prosumers with the ability to specify production and consumption offers with temporal flexibility. We solve the trading problem as a linear program, maximizing the energy traded over a long time horizon. We introduce a hybrid architecture for solving this linear program, combining the trustworthiness of distributed ledgers with the efficiency of conventional computational platforms. This hybrid architecture ensures the integrity of data and computational results—as long as majority of the ledger nodes are secure—while allowing the complex computation to be performed by a set of redundant and efficient solvers. The underlying communication substrate is implemented by using a distributed middleware, called Resilient Information Architecture Platform for Smart Grid (RIAPS) [13, 15].

Outline: We describe our futures trading approach, including extensions to support safety and privacy, in Section 2. We introduce the components of TRANSAX in Section 3, and detail its protocol in Section 4. We analyze how TRANSAX meets key requirements and discuss the tradeoff between privacy and efficiency in Section 5. We present an integrated testbed using GridLAB-D [8] and numerical results in Section 6. Finally, we give an overview of related work in Section 7, followed by conclusions in Section 8.

2 THE ENERGY TRADING APPROACH

In this section, we formalize our approach to *efficient energy trading* in TRANSAX. We then extend the formalization to support safety, privacy, and resilience.

We assume the distribution network infrastructure to be a microgrid with a set of feeders. A feeder has a fixed set of nodes, each representing a prosumer, which is a combination of load and distributed energy resources, such as rooftop solar panels and batteries. We assume that the prosumers can predict their future production and consumption based on historical data and anticipated utilization. The prosumers submit energy offers based on their predictions via automated agents that act on behalf of residents (i.e., residents do not need to trade manually). The predictions do not need to be perfect because we assume the existence of a distribution system operator (DSO), which also participates in the market and can supply residual demand not met through the local market. The DSO may use the market to incentivize timed energy production within the microgrid to aid in grid stabilization and in the promotion of related ancillary services [11] through updates to the price policy. To measure the prosumers’ energy production and consumption, a smart meter must be deployed at each prosumer. In practice, these smart meters must be tamper resistant to prevent prosumers from “stealing electricity”. After a smart meter has measured the net amount of energy consumed by the prosumer in some time interval, it can send the relevant information to the DSO for billing purposes.

Our goal is to find an optimal match between energy production and consumption offers, which we refer as the *energy trading problem*. Each offer is associated with an identity that belongs to the prosumer that posted the offer. We refer to these identities as *accounts*, and prosumers may generate any number of them.

2.1 Formalization

Let \mathcal{F} denote the set of feeders. On each feeder, there is a set of prosumers, who can make offers to buy and sell energy. We assume that time is divided into intervals of fixed length Δ , and we refer to the t -th interval simply as time interval t . For a list of symbols used in the paper, see Table 1.

For feeder $f \in \mathcal{F}$, we let \mathcal{S}_f and \mathcal{B}_f denote the set of selling and buying offers posted by prosumers in feeder f , respectively.³ A selling offer $s \in \mathcal{S}_f$ is a tuple (A_s, E_s, I_s, R_s) , where A_s is the account that posted the offer, E_s is the amount of energy to be sold, I_s is the set of time intervals in

³ To include the DSO in the formulation, we assign it to a “dummy” feeder.

Table 1. List of Symbols

Symbol	Description
Microgrid	
\mathcal{F}, \mathcal{U}	set of feeders and prosumers, resp.
C_f^{ext}, C_f^{int}	maximum allowed net and total, resp., power consumption or production in feeder $f \in \mathcal{F}$
EPL_u, ECL_u	maximum allowed production and consumption, resp., of prosumer u
C_g^{ext}, C_g^{int}	maximum allowed net and total, resp., power consumption or production in group $g \in \mathcal{G}$
EPA, ECA	asset granting permission to produce or consume, resp., a unit of energy
Δ	length of each time interval
T_{clear}	minimum number of time intervals between the finalization and notification of a trade
E_u^t	energy transferred by prosumer u in interval t
t_f	next interval to be finalized $t + 1 + T_{clear}$
Offers	
$\mathcal{S}_f, \mathcal{B}_f$	set of selling and buying offers, resp., from feeder $f \in \mathcal{F}$
\mathcal{S}, \mathcal{B}	set of all selling and buying offers, resp.
$\mathcal{S}^{(t)}, \mathcal{B}^{(t)}$	set of all selling and buying offers, resp., submitted by the end of time interval t
A_s, A_b	account that posted offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$, resp.
E_s, E_b	amount of energy to be sold or bought, resp., by offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$
I_s, I_b	time intervals in which energy could be provided or consumed by offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$, resp.
R_s, R_b	reservation prices of offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$, resp.
$\mathcal{M}(s), \mathcal{M}(b)$	set of offers that are matchable with offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$, resp.
Solution	
$p_{s,b,t}$	amount of energy that should be provided by s to b in interval t
$\pi_{s,b,t}$	unit price for the energy provided by s to b in interval t
$Feasible(\mathcal{S}, \mathcal{B})$	set of feasible solutions given sets of selling and buying offers \mathcal{S} and \mathcal{B}
$\hat{p}_{s,b,t}, \hat{\pi}_{s,b,t}$	finalized trade values
Implementation Parameters	
T_h	solve horizon (offers beyond horizon $t_f + T_h$ are not considered by solver)
$\hat{\Delta}$	length of the time step used for simulating the real interval of length Δ

which the energy could be provided, R_s is the reservation price, i.e., lowest unit price for which the prosumer is willing to sell energy. Similarly, a buying offer $b \in \mathcal{B}_f$ is a tuple (A_b, E_b, I_b, R_b) , where the values pertain to consuming/buying energy instead of producing/selling, and R_b is the highest price that the prosumer is willing to pay. For convenience, we also let \mathcal{S} and \mathcal{B} denote the set of all buying and selling offers (i.e., we let $\mathcal{S} = \cup_{f \in \mathcal{F}} \mathcal{S}_f$ and $\mathcal{B} = \cup_{f \in \mathcal{F}} \mathcal{B}_f$).

We say that a pair of selling and buying offers $s \in \mathcal{S}$ and $b \in \mathcal{B}$ is *matchable* if

$$R_s \leq R_b \text{ and } I_s \cap I_b \neq \emptyset. \quad (1)$$

In other words, a pair of offers is matchable if there exists a price that both prosumers would accept and a time interval in which the seller and buyer could provide and consume energy. For a given selling offer $s \in \mathcal{S}$, we let the set of buying offers that are matchable with s be denoted by $\mathcal{M}(s)$. Similarly, we let the set of selling offers that are matchable with a buying offer b be denoted by $\mathcal{M}(b)$.

A solution to the energy trading problem is a pair of vectors $(\mathbf{p}, \boldsymbol{\pi})$, where $p_{s,b,t}$ is a non-negative amount of energy that should be provided by offer $s \in \mathcal{S}$ and consumed by offer $b \in \mathcal{M}(s)$ in time interval $t \in I_s \cap I_b$ ⁴; and $\pi_{s,b,t}$ is the unit price for the energy provided by offer $s \in \mathcal{S}$ to offer $b \in \mathcal{M}(s)$ in time interval $t \in I_s \cap I_b$.

A pair of vectors $(\mathbf{p}, \boldsymbol{\pi})$ is a feasible solution to the energy trading problem if it satisfies the following two constraints. First, the amount of energy sold or bought from each offer is at most the

⁴ We require the both seller and buyer to produce a constant level of power during the time interval. This can be achieved by smart inverters.

amount of energy offered:

$$\forall s \in \mathcal{S} : \sum_{b \in \mathcal{M}(s)} \sum_{t \in I(s,b)} p_{s,b,t} \cdot \Delta \leq E_s \quad \text{and} \quad \forall b \in \mathcal{B} : \sum_{s \in \mathcal{M}(b)} \sum_{t \in I(s,b)} p_{s,b,t} \cdot \Delta \leq E_b \quad (2)$$

Second, the unit prices are between the reservation prices of the seller and buyer:

$$\forall s \in \mathcal{S}, b \in \mathcal{M}(s), t \in I(s,b) : R_s \leq \pi_{s,b,t} \leq R_b \quad (3)$$

The objective of the energy trading problem is to maximize the amount of energy traded. The rationale behind this objective is maximizing the load reduction on the bulk power grid. Formally, an optimal solution to the energy trading problem is

$$\max_{(\mathbf{p}, \boldsymbol{\pi}) \in \text{Feasible}(\mathcal{S}, \mathcal{B})} \sum_{s \in \mathcal{S}} \sum_{b \in \mathcal{M}(s)} \sum_{t \in I(s,b)} p_{s,b,t} \quad (4)$$

where $\text{Feasible}(\mathcal{S}, \mathcal{B})$ is the set of feasible solutions given selling and buying offers \mathcal{S} and \mathcal{B} (i.e., set of solutions satisfying Equations (2) and (3) with \mathcal{S} and \mathcal{B}).

The above formulation ensures feasibility, which takes the reservation prices into account. However, we do not address how to set the clearing prices in this paper. Clearing prices could be set using an existing approach, e.g., double auction; we discuss how this could be done in Section 8.1.

2.2 Safety Extensions

To ensure the safety of the microgrid, we introduce additional constraints on the solution to the energy trading problem. Each prosumer u has independent production and consumption limits, which are denoted by ECL_u and EPL_u , respectively. Further, each feeder $f \in \mathcal{F}$, has a transformer for incoming power, which has a capacity rating. We let C_f^{ext} denote the capacity of the transformer of feeder f . Similarly, the distribution lines and transformers for transferring power within the feeder have capacity ratings as well. We let C_f^{int} denote the maximum amount of power that is allowed to be consumed or produced within the feeder at any moment.⁵ These constraints are physically enforced by the over-current relays of the circuit breakers and feeders.

Now we generalize and introduce the notion of groups. We note that groups can correspond to feeders and support the constraints that we introduced in the previous paragraphs. They allow us to support physical layouts other than strictly feeders, and it will be useful for privacy later. We define a *group* g to be a set of feeders (i.e., $g \subseteq \mathcal{F}$). We let \mathcal{G} be the set of all groups, and for each group $g \in \mathcal{G}$, we introduce group safety limits C_g^{int} and C_g^{ext} , which are analogous to feeder limits. A solution is safe if it satisfies the following three constraints. First, the amount of power flowing into or out of a prosumer is within the production and consumption limits in all time intervals:

$$\forall u \in \mathcal{U}, t : \sum_{s \in \mathcal{S}_u} \sum_{b \in \mathcal{B}} p_{s,b,t} \leq EPL_u \quad \text{and} \quad \forall u \in \mathcal{U}, t : \sum_{b \in \mathcal{B}_u} \sum_{s \in \mathcal{S}} p_{s,b,t} \leq ECL_u \quad (5)$$

where \mathcal{S}_u and \mathcal{B}_u are the sets buying and selling offers posted by accounts owned by prosumer u .

Second, the amount of energy consumed and produced within each group is below the safety limit in all time intervals:

$$\forall g \in \mathcal{G}, t : \max \left\{ \sum_{b \in \mathcal{B}_g} \sum_{s \in \mathcal{S}} p_{s,b,t}, \sum_{s \in \mathcal{S}_g} \sum_{b \in \mathcal{B}} p_{s,b,t} \right\} \leq C_g^{int} \quad (6)$$

This means that the sum of all the buying and the selling trades cannot exceed the safety limit.

⁵ In other words, limit C_f^{ext} is imposed on the net production and net consumption of all prosumers in feeder f , while limit C_f^{int} is imposed on the total production and consumption of prosumers in feeder f .

Third, the amount of power flowing into or out of each group is within the safety limit in all time intervals:

$$\forall g \in \mathcal{G}, t : -C_g^{ext} \leq \left(\sum_{s \in \mathcal{S}_g} \sum_{b \in \mathcal{B}} p_{s,b,t} \right) - \left(\sum_{b \in \mathcal{B}_g} \sum_{s \in \mathcal{S}} p_{s,b,t} \right) \leq C_g^{ext} \quad (7)$$

Of the expression in the middle of the inequality, the first term is the total amount of energy sold by the group, and the second term is the total amount of energy bought by the group. Any trades that occurred within the group, or could have, do not end up counting towards the C_g^{ext} limit. Note that expression compared to C_g^{int} in Equation (6) is the largest of the buying and selling matches, whereas the expression compared to C_g^{ext} in Equation (7) is the difference between the buying and selling matches; this means that the value compared against C_g^{int} will always be larger than the value compared against C_g^{ext} . Thus if $C_g^{int} < C_g^{ext}$ then the external limit will never trip, and so we only need to consider $C_g^{ext} \leq C_g^{int}$. Verifying if a solution is safe in this case is simple, the offers in the solution need only be associated with the group they came from and checked against the constraints.

2.3 Privacy Extensions

To protect prosumers' privacy, we let them use anonymous accounts when posting offers. By generating new anonymous accounts, a prosumer can prevent others from linking the anonymous accounts to its actual identity, thereby keeping its trading activities private. However, anonymous accounts pose a threat to safety. Since the energy trading formalization with safety extension (see Equations (5) - (7)) discussed earlier requires the offers to be associated with the prosumer to enforce prosumer-level constraints and with the group from which they originated in order to be able to enforce group-level safety constraints. Without these associations prosumers, can generate any number of anonymous accounts, and post selling and buying offers for large amounts of energy without any intention of delivering and without facing any repercussions. A malicious or faulty prosumer could easily destabilize the grid with this form of reckless trading. Consequently, the amount of energy that may be traded by anonymous accounts belonging to a prosumer must be limited.

To enforce the prosumer-level constraints we introduce the concept of energy production and consumption assets, which allows us to disassociate the limiting of assets from the anonymity of offers. First, an *energy production asset* (EPA) is tuple $(E_{EPA}, I_{EPA}, G_{EPA})$, where

- E_{EPA} is the permission to sell a specific non-negative amount of energy to be produced,
- I_{EPA} is the set of intervals for which the asset is valid, and
- G_{EPA} is the group that the asset is associated with.

Second an *energy consumption asset* (ECA) represents a permission to buy a specific amount of energy and is defined by the same fields. For this asset, however, the fields define energy consumption instead of production. Each prosumer u is only permitted to withdraw assets up to the limits EPL_u and ECL_u into a non-anonymous account.

These assets can be moved to anonymous accounts in an untraceable way, but still retain their group association and the sum of the assets remains constant. Production assets are required to post a selling offer, and consumption assets are required to post a buying offer. For the offer to be valid, the account posting the offer must have assets that cover the amount and intervals offered. When a trade is finalized the assets are exchanged. We will provide more details on how they fit into the trading approach in Section 4.

To enforce group-level safety we only provide group-level anonymity, meaning that an offer can be traced back to its group of origin, but not to the individual prosumer within the group. When

forming a group the safety constraints will need to be set appropriately. We will discuss how they should be set and the associated energy trading capacity costs in Section 5.2.

2.4 Iterative Solutions Extension

In our basic problem formulation, we assumed that all buying and selling offers \mathcal{B} and \mathcal{S} are available at once, and we cleared the market in one take. In practice, however, the market conditions and the physical state of the DSO and prosumers may change over time, making it advantageous to submit new offers. Updating or cancelling offers could also be beneficial, however we do not provide that functionality here, but we describe how it could be accomplished in Section 8.1. As new offers are posted we need to recompute the solution. While new offers can increase the amount of energy traded, the *trade values* $p_{s,b,t}$ and $\pi_{s,b,t}$ need to be *finalized* at some point in time. At the very latest, values for interval t need to be finalized by the end of interval $t - 1$; otherwise, participants would have no chance of actually delivering the trade.

Here, we extend the energy trading problem to accommodate a time-varying offer set (where offers can be unmatched, matched and pending, or matched and finalized), and a time constraint for finalizing trades. Our approach finalizes only trades that need to be finalized, which maximizes efficiency while providing safety. We assume that all trades for time interval t' (i.e., all values $p_{s,b,t'}$ and $\pi_{s,b,t'}$) must be finalized and the trading prosumers must be notified by the end of time interval $t' - T_{clear} - 1$ (see Fig. 1), where T_{clear} is a positive integer constant that is set by the DSO. In other words, if the current interval is t , then all intervals up to $t + T_{clear}$ have already been finalized. Preventing “last-minute” changes can be crucial for safety and fairness since it allows both the DSO and the prosumers to prepare for delivering (or consuming) the right amount of power. In practice, the value of T_{clear} must be chosen accounting for both physical constraints (e.g., how long it takes to turn on a generator) and communication delay (e.g., some participants might learn of a trade with delay due to network disruptions).

We let $\hat{p}_{s,b,t}$ and $\hat{\pi}_{s,b,t}$ denote the finalized trade values, and we let $\mathcal{B}^{(t)}$ and $\mathcal{S}^{(t)}$ denote the set of buying and selling offers that participants have submitted by the end of time interval t . Then, the system takes the following steps at the end of each time interval t . First, find an optimal solution $(\mathbf{p}^*, \boldsymbol{\pi}^*)$ to the extended energy trading problem:

$$\max_{(\mathbf{p}, \boldsymbol{\pi}) \in \text{Feasible}(\mathcal{S}^{(t)}, \mathcal{B}^{(t)})} \sum_{s \in \mathcal{S}^{(t)}} \sum_{b \in \mathcal{M}(s)} \sum_{\tau \in I(s,b)} p_{s,b,\tau} \quad (8)$$

subject to

$$\forall \tau \leq t_f : p_{s,b,\tau} = \hat{p}_{s,b,\tau} \quad (9)$$

$$\pi_{s,b,\tau} = \hat{\pi}_{s,b,\tau} \quad (10)$$

Second, finalize trade values for time interval t_f based on the optimal solution $(\mathbf{p}^*, \boldsymbol{\pi}^*)$:

$$\hat{p}_{s,b,t_f} := p_{s,b,t_f}^* \quad (11)$$

$$\hat{\pi}_{s,b,t_f} := \pi_{s,b,t_f}^* \quad (12)$$

By taking the above steps at the end of each time interval, trades are always cleared based on as much information as possible (i.e., considering as many offers as possible)⁶ without violating any safety or timing constraints. Note that here $\text{Feasible}(\mathcal{S}, \mathcal{B})$ now also includes the safety constraints (5), (6), and (7).

⁶ This includes offers for intervals beyond the finalization interval. Effectively, matches for an interval beyond finalization can be changed if a better solution is found; however, finalized matches are permanent and never changed.

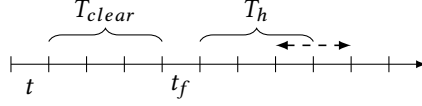


Fig. 1. Temporal parameters (t is current interval, t_f is the interval to be finalized).

2.5 Linear-Programming Solution:

To find the optimal solution efficiently, we frame the energy trading problem as a linear program. First, we create real-valued variables $p_{s,b,t}$ and $\pi_{s,b,t}$ for each $s \in \mathcal{S}, b \in \mathcal{M}(s), t \in I_s \cap I_b$. Then, the following reformulation of the matching problem is a linear program:

$$\max_{\mathbf{p}, \boldsymbol{\pi}} \sum_{s \in \mathcal{S}} \sum_{b \in \mathcal{M}(s)} \sum_{t \in I(s,b)} p_{s,b,t} \quad (13)$$

subject to the constraint Equations, which can all be expressed as linear inequalities (2), (3), (5), (6), (7), and

$$\mathbf{p} \geq \mathbf{0} \text{ and } \boldsymbol{\pi} \geq \mathbf{0}. \quad (14)$$

2.5.1 Practical Considerations. Even though Equation (4) can be formulated as a linear program and be solved efficiently (i.e., in polynomial time), the number of variables $\{p_{s,b,t}\}$ may grow prohibitively high as the number of offers and time intervals that they span increases. In practice, this may pose a significant challenge for solving the energy trading problem for larger transactive microgrids. A key observation that helps us tackle this challenge is that even though prosumers may post offers whose latest intervals are far in the future, the optimal solution for the finalized interval typically depends on only a few intervals ahead of the finalization deadline. Indeed, we have observed that considering intervals in the far future has little effect on the optimal solution for the interval that is to be finalized next (see Fig. 7).

Consequently, for practical solvers, we introduce a planning horizon T_h (see Fig. 1) that limits the intervals that need to be considered for a solution: for any $\hat{t} > t_f + T_h$, we set $p_{s,b,\hat{t}} = 0$, where t_f is the earliest interval that has not been finalized. By “pruning” the set of free variables, we can significantly improve the performance of the solver with negligible effect on solution quality (see Fig. 7). This results in the following “pruned” objective function:

$$\max_{(\mathbf{p}, \boldsymbol{\pi}) \in \text{Feasible}(\mathcal{S}, \mathcal{B})} \sum_{s \in \mathcal{S}} \sum_{b \in \mathcal{M}(s)} \sum_{\tau \in I_s \cap I_b \cap \{\tau; \tau \leq t_f + T_h\}} p_{s,b,\tau} \quad (15)$$

3 TRANSAX COMPONENTS

In this section, we introduce the components of TRANSAX (see Fig. 2), which implement the functionality described in the previous section. The components are built on top of RIAPS [15, 52]. For brevity, we do not elaborate on the role of RIAPS in this paper, and refer interested readers to [27].

3.1 DSO

The DSO is a trusted entity that manages the microgrid, handling financial operations like sending monthly bills, and functional operations like registering new smart meters.

Monthly bills are handled in conjunction with the smart meter. The DSO meets the prosumers’ residual demand and supply, i.e., consumption and production that they did not trade in advance due to inaccurate predictions or lack of trade partners. For each interval, the DSO sets the prices

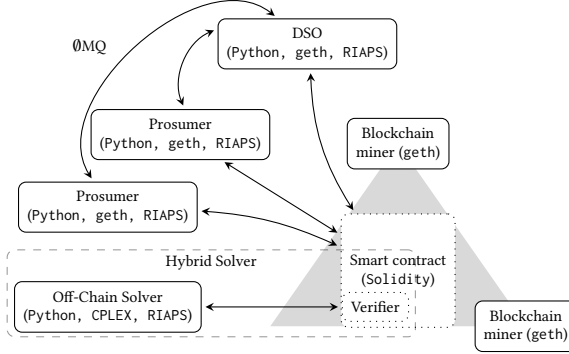


Fig. 2. Components of the energy trading system. In our reference implementation, we use Ethereum as the decentralized computation platform for smart contracts, and the other components interact with the blockchain network using the geth Ethereum client. The smart contract is implemented in Solidity, a high-level language for Ethereum, and it is executed by a private network of geth mining nodes. The off-chain solver uses CPLEX.

π_t^S and π_t^B for this residual energy consumption and production, respectively, which will be used by the smart meter for billing (see Section 4.4).

Registration means adding a new smart meter to the platform when it is installed for a prosumer. Registration includes assigning the smart meter to the prosumer’s group and setting the prosumer’s limits *EPL* and *ECL* according to the amount of energy the hardware connected to the smart meter can feasibly produce and consume respectively. The DSO can update these limits as required due to changes in the hardware monitored by the smart meter.

3.2 Prosumers

Prosumers are able to withdraw assets for each interval from the platform up to the limits set by the DSO, and trade those assets with other prosumers. To protect their privacy, they may also create new anonymous accounts and use them for trading. To maintain the anonymity of these accounts, prosumers always transfer assets to them through a *anonymizing mixer*. The mixer ensures that accounts cannot be linked to the prosumer that owns them. The mixer is implemented as a protocol that happens between prosumers, not through the DSO or other central entity. We discuss mixing and how it fits in the workflow in Section 4.2.

3.3 Smart Meter

The role of the smart meter is to measure the prosumer’s energy production and consumption, to monitor the assets allocated to the prosumer, and to provide privacy-preserving billing.

The smart meter checks that all assets withdrawn by a prosumer are accounted for. It also measures how much energy the prosumer transferred, ensuring that it does not exceed the *ECL* or *EPL* limits, and how much was not part of a trade. The amount that was not accounted for in trades is used to compute how much the prosumer owes the DSO for the interval. Each billing cycle (e.g., monthly) the smart meter sends that information to the DSO so it can bill the prosumer. This way the DSO has no fine-grained information on the energy profile of the consumer, it only knows the amount it needs to be paid for the energy it provided during that cycle. We discuss the workflow of the smart meter in Section 4.4.

3.4 Smart Contract

Smart-contracts are programs that provide trustworthy, general-purpose computation. Since they are trusted, they are suitable for financial settlements and have been implemented on various platforms including distributed ledgers. Contracts that utilize distributed ledgers are beneficial because they enable the decentralized enforcement of the contract rather than relying on a single trusted entity. We use a blockchain-based smart contract, since it provides a distributed ledger, to enable sensitive market features like asset trading and financial transactions.

The distributed ledger provides immutable storage service for offers, solutions, safety constraints, and a notification log of events. Prosumers and solvers check for these events and perform actions based on them. Offers are posted to the ledger using the smart contract functions. Valid offers must simply contain the format specified in Section 2 and the account posting the offer must have assets to cover what is offered. If an offer is valid, it is accepted and made available to other services through the smart contract.

3.5 Hybrid Solver

Although solving linear programs is not computationally hard, it can be challenging with a large number of variables and constraints in resource-constrained computing environments. Since computation is relatively expensive on blockchain-based distributed platforms⁷, solving even the “pruned” energy trading problem from Equation (15) might be infeasible using a blockchain-based smart contract. In light of this, we propose a *hybrid implementation approach*, which combines the trustworthiness of blockchain-based smart contracts with the efficiency of more traditional computational platforms.

The key idea of our hybrid approach is to (1) use high-performance computers to solve the computationally expensive linear program *off-blockchain* and then (2) use a smart contract to record the solution *on the blockchain*. To implement this hybrid approach securely and reliably, we must account for the following issues.

- Computation that is performed off-blockchain does not satisfy the auditability and security requirements that smart contracts do. Thus, the results of any off-blockchain computation must be verified in some way by the smart contract before recording them on the blockchain.
- Due to network disruptions and other errors (including deliberate denial-of-service attacks), the off-blockchain solver might fail to provide the smart contract with a solution on time (i.e., before trades are supposed to be finalized). Thus, the smart contract must be able to proceed without an up-to-date solution.
- For the sake of reliability, the smart contract should accept solutions from multiple off-blockchain sources; however, these sources might provide different solutions. Thus, the smart contract must be able to choose from multiple solutions (some of which may come from a compromised computer).

3.5.1 Smart Contract based Solution Verification. Apart from storing data (Section 3.4), our smart contract is also designed to (1) verify whether a solution $(\mathbf{p}, \boldsymbol{\pi})$ is feasible and (2) compute the value of the objective function for a feasible solution. Compared to finding an optimal solution, these operations are computationally inexpensive, and they can easily be performed on a blockchain-based decentralized platform. Using these capabilities, the smart contract provides the following functionality:

- Solutions may be submitted to the contract at any time. The contract verifies the feasibility of each submitted solution, and if the solution is feasible, then it computes the value of the objective

⁷ Further, Solidity, the preferred high-level language for Ethereum, currently lacks built-in support for certain features that would facilitate the implementation of a linear programming solver, such as floating-point arithmetics [54].

function. The contract always keeps track of the best feasible solution submitted so far, which we call the *candidate solution*.

- At the end of each time interval t , the contract finalizes the trade values for interval $t_f = t + T_{clear} + 1$ based on the candidate solution.⁸

3.5.2 Off-blockchain Matching Solver. We complement the smart contract with an efficient linear programming solver, specifically CPLEX⁹, which can be run off-blockchain, on any capable computer (or multiple computers for reliability). The solver is run periodically to find a solution to the energy trading problem based on the latest set of offers posted. Once a solution is found by the matching solver, it is submitted to the smart contract in a blockchain transaction. Note that if new offers have been posted since the solver started working on the solution, the contract will still consider the solution to be feasible. This is because any feasible solution for sets \mathcal{S} and \mathcal{B} also being feasible for supersets $\mathcal{S}' \supseteq \mathcal{S}$ and $\mathcal{B}' \supseteq \mathcal{B}$.

From the perspective of the solver, being able to submit multiple solutions to the contract for the same problem has many advantages. For example, it allows the linear programming solver to be run as an anytime algorithm. Further, we can allow multiple—potentially untrusted—entities to try to solve the problem and submit solutions, since the smart contract will always choose the best feasible one. This is especially important in microgrids where a trusted third party is not guaranteed to always be present. In such settings, prosumers can be allowed to volunteer and provide solutions to the energy trading problem.¹⁰

Thereby, we enable finding solutions in an efficient and very flexible manner, while reaping the benefits of smart contracts, such as auditability and trustworthiness.

4 TRANSAX PROTOCOL

In this section, we describe how the components introduced in the previous section interact via the TRANSAX protocol, which is depicted in Fig. 3

4.1 Registration

When a new customer is added to the grid, a smart meter is installed. The DSO registers the smart meter by calling ①¹¹ *registerSmartMeter* on the TRANSAX smart contract. This call sets the asset allocation limits for that customer and records which feeder it is located on in the grid. The customer then registers as a prosumer with TRANSAX by calling ② *registerProsumer*.

The registration information requires each prosumer to specify a smart meter, and to provide a DSO certified public address that corresponds to the specified smart meter for the DSO to use when allocating assets. Since the smart meter is associated with a specific feeder, the smart contract adds the prosumer to the group associated with that feeder. This is required to ensure that feeder-level safety constraints can be correctly applied. The registrations can happen asynchronously, allowing new prosumers to join at any time, even long after trading has commenced. The registration process occurs only once for each smart meter and prosumer. Once registered, a prosumer may participate in the following trading protocol repeatedly.

⁸ If no solution has been submitted to the contract so far, which might be the case right after the trading system has been launched, $\mathbf{p} = \mathbf{0}$ may be used as a candidate solution. ⁹ <https://www.ibm.com/analytics/cplex-optimizer> ¹⁰ Of course, each prosumer will try to submit a solution that favors the prosumer. However, the submitted solution still needs to be superior with respect to the optimization objective, which roughly corresponds to social utility. Hence, each prosumer is incentivized to improve social utility by submitting superior solutions that favor the prosumer. We leave the analysis of these incentives for future work. ¹¹ The circled numbers correspond to the numbered edges in Fig. 3

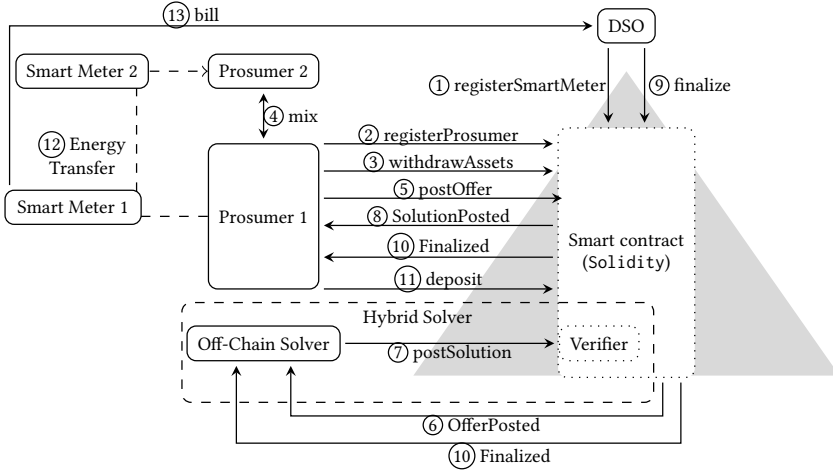


Fig. 3. Example workflow of TRANSAX. Nodes represent entities in the platform, and edges represent interactions, such as smart-contract function calls. In this example, prosumer 1 is selling energy to prosumer 2 and the dashed line represents the energy transfer.

4.2 Mixing

Once a prosumer has registered, it may withdraw assets into their public address for future intervals by calling **3** *withdrawAssets*. After withdrawing assets, a prosumer could make offers using *postOffer*. However, if it made offers using its public account, then the trades could be traced back to the prosumer, thereby violating the privacy requirements. Instead, the prosumer creates anonymous addresses and transfers the assets from its public address to the anonymous addresses via **4** mixing assets with other prosumers. They can do this within their assigned groups by executing a decentralized mixing protocol such as the one described in [46]. Within the group, each prosumer supplies anonymous addresses and the amount to be deposited in each address. They also supply a public address containing the assets that will be transferred to the anonymous addresses. Then, the prosumers execute the protocol, transferring the assets anonymously, without associating them with the public addresses. If the prosumer transferred assets to a new “anonymous” account directly without using mixing, then the new account would not truly be anonymous as there would be no difficulty in tying it back to the prosumer.

4.3 Trading

Now the prosumers can post anonymous offers using their anonymous accounts by calling function **5** *postOffer*. The smart contract checks that the account posting the offer has assets that cover the amount and intervals specified in the offer. If not then the offer is rejected. If the offer is valid the smart contract then emits event **6** *OfferPosted*, notifying the off-chain matching solvers. The matching solvers may wait for many prosumers to post many offers, but eventually, it pairs buying and selling offers and posts the solutions by calling function **7** *postSolution*. The smart contract checks the solution to make sure that it is feasible according to the feasibility requirements described in Section 2, including checking that the trades do not exceed the group capacity constraint. If the solution is valid, then smart contract saves it and emits event **8** *SolutionPosted*, notifying the prosumers of the current candidate solution. Additional solutions may be submitted by any solver,

and if those solutions are valid and superior, *i.e.*, they trade more energy, then the smart contract will update the candidate solution.

4.4 Billing

As an interval comes to a close, the DSO calls¹² function ⑨ *finalize* which means that offers for interval t_f are no longer accepted and the smart contract transfers funds from the consuming offer's account to the producing offer's account. It also exchanges the *EPA* assets of the seller for the *ECA* assets of the buyer and vice versa for each of the matched offers. The call also emits the ⑩ *Finalized* event, notifying the solvers to update their solving interval, and the prosumers that the trades for interval t_f have been finalized. If a prosumer posts offers with many anonymous accounts, it will have to aggregate all the corresponding trades to determine how much power it is expected to produce/consume during that interval when it arrives. Once the prosumers are notified of the trades, they call function ⑪ *deposit* to transfer all assets for the finalized interval from the prosumers anonymous accounts to an anonymous account owned by their smart meter.

The smart meter checks that the total amount of assets deposited matches the amount withdrawn for the finalized interval. This ensures that there are no trades that have not been accounted for. It also compares the total of all production assets that were deposited against the production originally withdrawn to compute the net energy sold ($\Delta EPA = EPA_{deposit} - EPA_{withdraw}$). Then when the interval t arrives and the power is ⑫ transferred, the smart meter measures energy production/-consumption. Let E_u^t be net energy production (negative value represent net consumption) of prosumer u at time interval t . The smart meter compares the net energy production (consumption) against the *EPL* (*ECL*) of the prosumer to make sure it does not violate the safety constraint. It then computes the difference between the net energy sold and the net energy production to get the residual production (again, negative values are residual consumption). The residual production or consumption is multiplied by the DSO sell or buy prices, respectively, to compute what the prosumer owes the DSO each interval. Every ⑬ billing cycle the smart meter sums the cost of the residuals and sends that to the DSO so the DSO can send the monthly bill. The bill B_u^t of prosumer u for timeslot t , which will be paid by the prosumer to the DSO, is

$$B_u^t = \begin{cases} (E_u^t + \Delta EPA) \cdot \pi_t^S & \text{if } E_u^t + \Delta EPA < 0 \\ (E_u^t + \Delta EPA) \cdot \pi_t^B & \text{otherwise,} \end{cases} \quad (16)$$

where π_t^S is how much the DSO pays to purchase power and π_t^B is how much the DSO charges for power. The prices could be functions of $E_u^t + \Delta EPA$, to charge more as the deviation from the predicted energy requirements increase. The price schedule is set for each timeslot t by the DSO.

5 DISCUSSION AND ANALYSIS

In this section, we first describe how the TRANSAX design ensures the security, resilience and safety of the system. Then, we provide a discussion on the inherent trade-offs between efficiency, and privacy.

5.1 Requirement Evaluation

5.1.1 Security and Safety. The underlying blockchain platform provides basic security features, so we are not concerned with the operations occurring on the blockchain. We are concerned with the secure and reliable operation of the solver. Similarly, the basic safety of the system is handled by the constraints described in Section 2.2. The safety constraints are applied correctly and reliably by

¹² Note that by default the DSO calls the *finalize* function to increment the current interval, but since this function is time guarded, any other entity can call it, which provides additional resilience.

the same contract. An adversary cannot force the contract to finalize trades based on an unsafe (i.e., infeasible) solution since such a solution would be rejected. Similarly, an adversary cannot force the contract to choose an inferior solution instead of a superior one. In sum, the only action available to the adversary is proposing a superior feasible solution, which would actually improve energy trading in the microgrid.

5.1.2 Resilience. Now we show that our contract is reliable and can tolerate temporary disruptions in the DSO, solvers, or the communication network. First, since the *finalize* contract function is time guarded any entity can call it, and the system can progress without a DSO which is only required for registering new prosumers and their smart meters. Second, notice that any solution $(\mathbf{p}, \boldsymbol{\pi})$ that is feasible for sets \mathcal{S} and \mathcal{B} is also feasible for supersets $\mathcal{S}' \supseteq \mathcal{S}$ and $\mathcal{B}' \supseteq \mathcal{B}$. As the sets of offers can only grow over time, the contract can use a candidate solution submitted during time interval t to finalize trades in any subsequent time interval $\tau > t$. In fact, without receiving new solutions, the difference between the amount of finalized trades and the optimum will increase only gradually: since the earlier candidate solution can specify trades for any future time interval, the difference is only due to the offers that have been posted since the solution was found and submitted. Thus, the system can continue making trades using older valid solutions

5.1.3 Trading Efficiency. The trading platform we have presented is able to support efficient trading through temporal flexibility. We show this through Example 1. As a reminder, this is due to prosumers being able to specify their production/consumption capacities and preferences (i.e., reservation prices) via offers and the linear-program finding an optimal matching. In Section 6.2, we show using simulation that energy trading reduces the load on the power grid.

Example 1. Consider a community with two prosumers (P_1, P_2) and one consumer (C_1) on a single feeder. We divide each day into 15 minute intervals. Let us assume that P_1 has the ability to transfer 10 kW into the feeder during interval 48, which translates to 12:00pm–12:15pm. Assume similarly that P_2 can also provide 30 kW to the feeder in interval 48, but it has battery storage. Since P_2 has battery—unlike P_1 , who must either transfer the energy or waste it— P_2 can delay the transfer until a future interval, e.g., interval 49. Now suppose that C_1 needs to consume 30 kW in interval 48 and 10 kW in interval 49. A possible solution would be to provide all 30 kW to C_1 from P_2 in interval 48. However, that will lead to the waste of energy provided by P_1 . Thus, a better solution will be to consume 10 kW from P_1 in interval 48 and 20 kW from P_2 in interval 48. Then, transfer 10 kW from P_2 in interval 49, which is more efficient than the first matching as it allows more energy (summed across the intervals) to be transferred. Thus, we see that permitting temporal flexibility can significantly increase trading volume, though it does increase the size of the optimization problem, increasing computational complexity.

5.1.4 Privacy. The platform provides pseudo-anonymity as the individual offers cannot be tied back to the prosumer who posted them since the offer is only affiliated with an anonymous address and contains only the energy amount and reservation price. Additionally the DSO does not know the total amount of energy utilized by the prosumers thanks to the anonymous billing via the smart meter. However, to preserve safety, some information about the prosumers needs to be public to allow checking of the offers to ensure that they are safe, or limit the resources available to them.

In our design, we assume that the consumption (*ECL*) and production (*EPL*) limits of each prosumer are public information, as well as which feeder a prosumer is on. The group safety constraints C_g^{int} and C_g^{ext} are also public. Recall that the smart contract ensures that no prosumer can withdraw more assets than the specified limits, and that any offer which violates the recorded safety constraints will be rejected. As a result, the only way to violate the safety requirements is if the asset limits or safety constraints are set incorrectly, which is not allowed by our design.

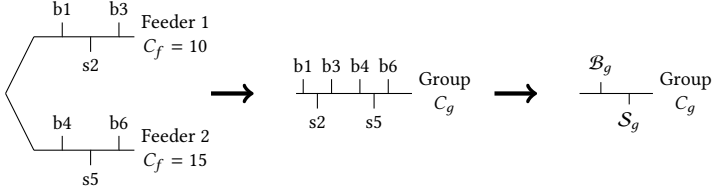


Fig. 4. Feeder conversion diagram.

However, as we will show below it is possible to improve privacy by choosing a conservative safety constraint for a group or a conservative limit on the maximum assets a prosumer can withdraw, which impacts the trading efficiency. Consider the following example for illustration.

Example 2. Given a group g with an internal constraint of $C_g^{int} = 30$ with prosumer $p1$ with $EPL = 10$ and $p2$ with $EPL = 20$. Assume that the prosumers in this group have anonymized their assets. If the total assets traded by the group—call it T_f —is below 10 there is no way to definitively say that either prosumer is trading. If the assets traded by f exceeds 10 then we know that $p2$ is trading at least $T_f - 10$ since $p1$ can only produce 10. If $T_f > 20$ then we know that $p1$ is trading at least $T_f - 20$. If $T_f = 30$ or 0 then we know the full state of the feeder, either both are at maximum capacity or are off. To improve anonymity, the feeder as a whole should not trade more than 10. This however reduces trading efficiency considerably. However if both prosumers have an $EPL = 15$ then anonymity is improved until trading exceeds 15. Thus it is important to select the constraints carefully. We discuss this in Section 5.2.

5.2 Tradeoff between Privacy and Efficiency

Note that safety of the system is a primary requirement and we cannot compromise it.

5.2.1 Selecting Group Constraints. When the group is limited to a single feeder, the safety constraints are simply the feeder's constraints. However, in a group, once the accounts are anonymous, we cannot tell which feeder they belong to. Thus, to preserve safety, the constraints need to be modified. To do this, the set of feeders are transformed into a group by treating all the prosumers in those feeders as though they were on a common feeder. Since the offers are anonymous at the group-level, the system can treat the group as a single feeder with two prosumers: one which posts production offers and one which post consumption offers (see Fig. 4).

Now we discuss how the group constraints can be set to ensure safety, and how much this costs in terms of trading efficiency. We define this as the cost of privacy:

Definition 5.1. Privacy Cost:

In order to safely provide privacy, some amount of energy transfer (which would otherwise be satisfied by the trades) is lost. This is the cost of privacy.

First, EPL_u and ECL_u are the same type of constraint representing a negative or positive flow of energy, so we will use EL_u to represent both in our analysis, but in each case the equations refer to both. We also have the smart contract set $C_g^{ext} = C_g^{int}$ and refer to it as C_g for now.

When setting the group safety limits, there are two cases to consider. We present them below.

Case 1. There is a set of prosumers in the group that is capable of exceeding the safety constraint of the feeder they are on.

Assume a microgrid with feeders \mathcal{F} and groups \mathcal{G} , wherein EL for each prosumer and C_f for each feeder can have any value. In order for this system to be safe, C_g for every group g must be:

$$C_g \leq \min \left\{ C_f \mid f \in g \text{ and } \sum_{u \in f} EL_u \geq C_f \right\} \quad (17)$$

PROOF. Assume Equation (17) is false, and that the system is safe. This means $\exists C_f < \sum_{u \in f} EL_u$ and $C_g > C_f$. Let $\sum_{u \in f} EL_u = C_g$ —then the prosumers in f can trade EL assets. This exceeds the feeder safety limit, and the system is not safe. Equation (17) must therefore be true. \square

Thus, the best value for the group constraint is when Equation (17) is equality. This means that the group as a whole can at most produce the same amount as the single smallest of its internal feeders. The cost in this case is:

$$\text{cost} = \min \left\{ \sum_{\forall s \in \mathcal{S}_g} E_s, \sum_{\forall b \in \mathcal{B}_g} E_b \right\} - \min \left\{ \sum_{\forall s \in \mathcal{S}_g} E_s, \sum_{\forall b \in \mathcal{B}_g} E_b, C_g \right\}. \quad (18)$$

Simply, the cost is the amount by which the potential trades exceed the safety constraint.

Case 2. No set of prosumers in any of the feeders in the group are capable of exceeding their feeders' safety constraint.

Given a microgrid with feeders \mathcal{F} and groups \mathcal{G} where C_f can have any value and

$$\forall_g \forall_{f \in g} \sum_{u \in f} EL_u \leq C_f, \quad (19)$$

group constraint should be set as $C_g = \sum_{f \in g} C_f$ to maximize trading, and trades can be done safely.

PROOF. Assume a microgrid is not safe and Equation (19) is true. Then, $\exists f$ such that $\sum_{u \in f} EL_u > C_f$. But, Equation (19) says this is not allowed. So, the system is safe. \square

In this case, there is no cost to group privacy. The safety is provided by the the asset withdrawal limits rather than the group constraint. Note that case 1 can be converted to case 2 by modifying the prosumer limits so that the prosumers on a feeder cannot access assets that exceed their feeder's safety constraint. How the production and consumption limits are set can be negotiated by the DSO and the prosumers during installation.

A consideration when converting from case 1 to case 2: since the prosumers have the capability to exceed C_f^{int} , they also exceed C_f^{ext} . Now, if instead of being equal (as we assumed in the beginning), $C_f^{int} > C_f^{ext}$, then there is a cost associated with adding that feeder to a group, where the maximum cost is $C_f^{int} - C_f^{ext}$.

This cost is incurred because trades internal to the feeder can't be distinguished from trades which are external, so the total amount traded by the feeder must not exceed C_f^{ext} . Otherwise, the external feeder over-current relay would trip. As a result, the prosumer must not be able to withdraw assets greater than their external feeder limit, resulting in the cost.

5.2.2 Insights on Grouping. Based on the analysis of the effects of privacy on efficiency, the best strategy is to limit the trading assets of the prosumers such that they remain less than the feeder constraints. This means that all feeders can be safely grouped. The cost of grouping feeders is the loss of flexibility in trading due to the rigid asset limits. The cost will be at most the feeder limit minus the prosumer asset limit, if that prosumer has the capacity to reach the feeder limit, and if no other prosumers in its feeder are trading. This could be mitigated by an additional mixing and

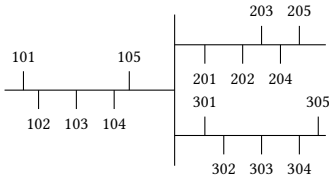


Fig. 5. Topology of the simulated distribution network.




Fig. 6. Message structure between simulator and TRANSAX.

trading step within the feeder, but we have not examined this possibility in detail. There is a second criterion that may influence grouping decisions. There is information leakage, and at the extremes (max load, zero load) anonymity ceases to exist. We assume that generally this will not be the case, and the odds of that occurring diminish if there are many feeders in the group. Information leakage can be reduced by setting all the asset limits to the same value for all prosumers. The maximum system cost of this is the difference between the feeder limit and the sum of the prosumer limits. To reduce information leakage, groups should consist of feeders with similar limits.

6 EXPERIMENTAL EVALUATION

In this section, we present a simulation testbed that we developed for evaluating TRANSAX, as well as our initial results illustrating its effectiveness in reducing the load on the bulk power grid.

The system to demonstrate the simulation platform has three major parts. First, the simulation is controlled by a Python agent that sends simulation status and receives actuation commands from TRANSAX. Second, we use GridLAB-D as a discrete-time distribution network simulator. Third, messages and time steps between the Python agent and GridLAB-D are coordinated by the Framework for Network Co-Simulation¹³.

6.1 Testbed

6.1.1 Distribution System. The distribution system is modeled using GridLAB-D and is simulated on an x86 computer with an i7 processor and 24GB of RAM. The distribution topology (see Fig. 5) consists of a single substation feeding three main overhead lines that are connected to prosumers. The lines below the main lines represent prosumers that are modeled as houses with battery and solar panels that can either consume or produce energy depending on the net output of the solar panels and battery, and those above represent prosumers that have loads and only consume. For the demonstration discussed in this paper, the simulation was built with 9 producer nodes and 6 consumer nodes. The control logic of each prosumer is executed on a single board computer (specifically the BeagleBone Black¹⁴).

6.1.2 Simulation Time Synchronization. The Python agent is the interface that relays data between the GridLAB-D simulation and TRANSAX, as well as synchronizes time between GridLAB-D's variable time-step solver and TRANSAX's matching solver. The solution frequency is a system parameter. In this demonstration, TRANSAX posted trades for each 15-minute interval of logical time.

The most important feature of this demonstration is its methodology for synchronizing time between GridLAB-D and TRANSAX. As noted above, GridLAB-D and TRANSAX are time synchronized through the Python Agent. The Python agent forces GridLAB-D's variable time step solver to pause at each logical 15-minute interval. The Python agent waits until it receives a step request to continue from TRANSAX. While GridLAB-D is paused, the Python agent and TRANSAX

¹³ <https://github.com/FNCS/fncs> ¹⁴ <https://beagleboard.org/black>

send all the necessary messages for TRANSAX to find and post finalized trades. After 2 minutes of real-time, TRANSAX sends a step request to the Python agent, and the Python agent actuates its control parameters for the next interval and commands GridLAB-D to simulate the next 15 minutes of logical time. This process recurs iteratively for the duration of the simulation’s logical time. The time synchronization strategy outlined is scalable to any desired time period for the TRANSAX solver. The strategy provides freedom to run experiments such as how the solver’s time period effects the energy traded, the stability of the finalized trades, and the computational complexity.

6.1.3 System Message Structure. The general message structure between GridLAB-D, the Python agent, and TRANSAX is shown in Fig. 6. While GridLAB-D is paused, TRANSAX agents request charge status for their batteries in the GridLAB-D simulation. They use this data, along with their predicted energy usage, to create a bid which is sent to TRANSAX. TRANSAX agents send the finalized trades back to the Python agent. The Python agent sets each simulated node’s power output for the next interval based on the finalized trades from TRANSAX by modifying GridLAB-D system parameters. In this demonstration, the Python agent only meets the finalized trades by modifying battery power outputs. However, the Python agent has control over all the dynamically modifiable parameters in GridLAB-D. As a consequence, future demonstrations could incorporate more control parameters such as curtailments to solar output or curtailments to energy used by pure consumers.

6.1.4 Setup. The experiment setup to test TRANSAX and the simulation contained 9 producers and 6 consumers for a total of 15 nodes in the distribution network. The simulation ran in logical time from 08:00 to 20:00 for a total of 12 hours. The experiment was run in two scenarios. First, the simulation was run without a link to TRANSAX. In this scenario, the batteries were not set to output any power to demonstrate the behavior of the system without a transactive energy system. Second, the simulation was run in conjunction with TRANSAX to demonstrate the efficacy of the transactive energy solver. During our experiments, we speed up the simulation by letting the real-time length of the time interval be $\hat{\Delta} < \Delta$ where $\hat{\Delta} = 2min$ and $\Delta = 15min$. Note that $\hat{\Delta}$ is the amount of real time passed in the simulation before proceeding to the next interval. This allows us to speed up the experiments without compromising our results since running the system slower would be easier.

6.2 Results

6.2.1 Varying the time horizon. We ran multiple simulations of a microgrid varying the value of T_h and selected a time interval during which we measured the memory usage and energy traded. In Fig. 7, we see that as the time horizon increases so does the memory usage and energy traded until $T_h = 30$ at which point there is not additional gain to energy traded. The time horizon also impacts the CPU utilization of the solver (not shown). This demonstrates that we can improve solver performance and still obtain quality solutions.

6.2.2 Trading Impact. The simulation was first run without battery power output and without any control by TRANSAX. This output was used to generate an energy profile for each prosumer for each interval. Then, the simulation was repeated with the prosumers submitting offers matching their energy profile to the TRANSAX system to test the effect that batteries can have to mitigate load and utilize solar overproduction. For this reason, this test represents an ideal scenario with accurate bids for each 15 minute interval. Fig. 9 shows the comparison of substation loads with and without TRANSAX. The horizontal axis is the simulated time since the start of the simulation. The vertical axis shows the power load on the substation, negative values signify that prosumer generation exceeds their loads. Without TRANSAX the system solar generation begins to outproduce the

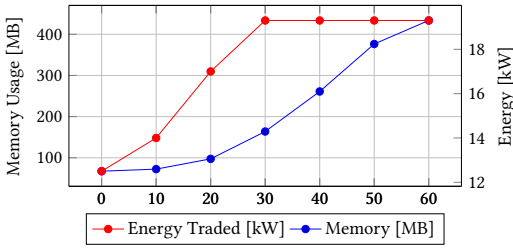


Fig. 7. Memory consumption and energy traded during a single interval of the simulation for various values of T_h using the CPLEX solver.

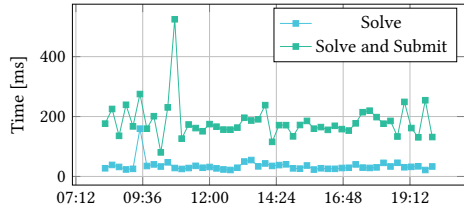


Fig. 8. *Solve* time is how long it took the solver to find a solution to the energy trading problem. *Solve and Submit* time is how long it took to find the solution and submit it to the smart contract.




Fig. 9. Substation load with and without TRANSAX.




Fig. 10. Average battery charge level.

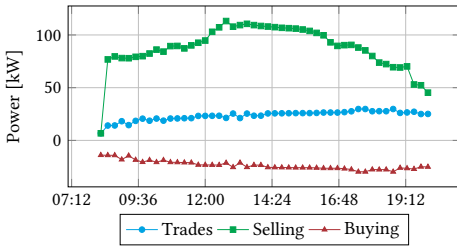


Fig. 11. Green: sum of all production offers for each interval. Red: negative sum of all consumption offers for each interval. Blue: sum of all energy traded in each interval, whose maximum value is the minimum of the production and consumption offers.

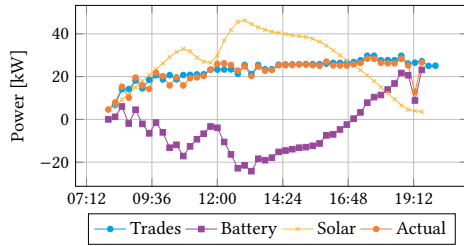


Fig. 12. Yellow: simulated solar profile. Purple: simulated battery charge level. Orange: simulated energy traded. Blue: total energy trades recorded in the market.

system load within the first interval. The solar reaches its peak production around 12:45. Finally, at 16:45, the load becomes greater than the solar production and the substation load is positive. The inclusion of TRANSAX dramatically reduces the need for the substation backup. From 8:00 to 16:45, the overproduction of solar meant that the batteries on the system were charging, which mitigated the negative load on the substation. After 16:45, the batteries on the system discharged and mitigated the positive load on the system. Fig. 10 shows the average battery charge level across all 9 batteries. At the end of the simulation (20:00), the average battery only had a 25% charge. This means that if the simulation were to go further into the night, there would not have been enough battery capacity to power the house loads for the entirety of the night.

Fig. 11 shows the total amount of energy offered for each interval, as well as the total amount of energy recorded in trades. In Fig. 12 we see that the trades recorded (Blue) are reasonably consistent

with the measured load (orange) on the system, with one notable exception at 14:15. The deviations occur because currently the prosumers assume that over an interval the solar input will remain constant, and this value is used when making offers.

Fig. 8 shows the time required by the platform to find the optimal matching of a set of offers (green), as well as that time combined with the time required to submit that solution to the smart contract (blue). The majority of the time spent is due to the smart contract communications.

The results of the simulation with TRANSAX are promising. TRANSAX found energy trade solutions for each interval that resulted in an overall mitigation of substation load. The distribution was not completely independent of the substation feeder; there was still a need for a connection to the larger distribution grid.

Finally, these results demonstrate that the developed simulation platform can effectively time synchronize and integrate with a transactive energy system like TRANSAX. This simulation platform is flexible and reusable, and it can be used in future experiments to thoroughly investigate the performance of TRANSAX for many different parameters.

7 RELATED WORK

7.1 Existing Deployments

Existing deployments are limited and so far have not published results. The closest is Wörner et al. [55], who have developed an implementation of their peer-to-peer energy market and deployed it to a town in Switzerland. Their goal is to gather empirical evidence to answer the question of what are the benefits of a blockchain system in the electricity use case. Their study will conclude at the end of 2019, at which time they will analyze the data to determine the performance of the blockchain system. To motivate the design of their system, they carried out a targeted literature review, and selected “a double auction with discriminative pricing as the most suitable market mechanism for electricity exchange.”

7.2 Buyer-Seller Matching

After prosumers presented their energy availabilities and demands in form of offers, these offers need to be matched. Researchers have proposed two approaches for this problem.

7.2.1 Stable Matching. Stable matching refers to matching of all possible buy and sell offers in a bipartite graph. Yucel *et al.* proposed a homomorphic encryption-based position hiding method [57] which protects users’ privacy from adversary matchers. Nunna *et al.* [40] proposed the symmetrical allocation problem based on native auction algorithm to match buyers and sellers. PowerLedger [32] uses another mechanism to match offers. Offers are broken into equal portions and matched together e.g., when a new consumer arrives, it will receive the equal allocation from the energy pool in the area.

7.2.2 Auction. Another approach to match buyers’ offers to sellers’ is to use auctioning approaches. Majumder *et al.* [33] proposed a double auction mechanism before the era of blockchains where the controller doesn’t need the users’ data. As a result, the privacy of bidders and sellers will be preserved. In the era of blockchains, Kang *et al.* [22] and Guerrero [18] used double auctions to match parties and not goods in blockchains. To ensure integrity of results of matching, Wang [53] proposed a multi-signed digital certificate. Khorsani *et al.* [24] designed a greedy algorithm with the averaging auction mechanism to match buyers with higher price to sellers with lower prices. Zhao *et al.* [61] created a two phase auctioning algorithm to find the optimal pricing for bids. Finally, Zhang *et al.* [59] developed a non-cooperative auctioning game and used it to find the optimal solution for the matching problem using the Nash Equilibrium.

7.3 Online Information Management Platforms

Information Management Platform are designed to collect, process, transmit, and analyze data. In this context, data collection usually happens at the edge because that is where edge devices with sensors are deployed to monitor surrounding environments. TRANSAX does not suggest a specific data collection methodology. Rather, it follows an actor-driven design pattern where “prosumer” actors can integrate their own agents into TRANSAX by using the provided APIs. Another concern of these platforms is the cost of processing. Traditionally, this problem was solved using scalable cloud resources in-house [47]. However, TRANSAX enables a decentralized ecosystem, where components of the platform can run directly on edge nodes, which is one of the reasons why we designed it to be asynchronous in nature.

To an extent, the information architecture of TRANSAX can be compared to dataflow engines, such as Storm¹⁵, Spark¹⁶, and S4 [38]. All of these existing dataflow engines use some form of a computation graph, comprising computation nodes and dataflow edges. These engines are designed for batch-processing and/or stream-processing high volumes of data in resource intensive nodes, and do not necessarily provide additional “platform services” like trust management or solver nodes.

7.4 Grid Control and Stability

One integral part of smart-grids are the microgrid controllers which ensure stability and resiliency of the microgrid. They enable transition of the microgrid from grid-connected to islanded [30, 50] so that the failures in the grid do not cascade to other areas similar to the outage event back in 1999 Sao Paulo, Brazil [60]. Currently, most of microgrid controllers are centralized [23] which are vulnerable to cyber-threats and privacy issues. A large spectrum of cyber-threats are applicable on centralized microgrid controllers with single-point-of-failure ranging from attackers eavesdropping on channels between the controllable resource and the centralized controller to steal critical information of the users or network infrastructure, performing DDOS attacks on the centralized controller, or manipulation of demand via IoT (MadIoT) attacks[19] to injecting malware into the market operation system and manipulate settings, such as DLMP limits or clearing time interval similar to the notable cyber-attack against Ukrainian power systems in December 2015 [28, 58].

Due to these drawbacks of centralized grid controls, industry is transforming from centralized to decentralized[48, 56]. The aim of TRANSAX is to create a decentralized transactive energy market which ensures privacy and security of users while maintaining stability and resiliency of the grid.

7.5 Security and Privacy

7.5.1 Communication Security. First step to preserve users’ privacy and anonymity in a distributed system is to provide communication privacy. Without this, an adversary can discern who is making a function call or sending a message over the network based on the sender’s MAC address, IP address, or route to destination. Existing protocols for low-latency communication anonymity include onion routing [43], the similar garlic routing [31], STAC [21], and the decentralized Matrix protocol¹⁷. However, Murdoch and Danezis [37] show that a low-cost traffic analysis is possible of the Tor-network, theoretically and experimentally. Communication security is an orthogonal research problem to TRANSAX.

7.5.2 Address Anonymity. Communication anonymity is necessary but not sufficient for anonymous trading, as the cryptographic objectives of authentication and legitimacy are not fulfilled. We suggest using cryptographic techniques from distributed ledgers, *blockchains*, and cryptocurrencies. The most adopted one, Bitcoin allows for very simple digital cash spending but has serious privacy and anonymity flaws [2, 4, 44]. Additionally, Biryukov and Pustogarov, 2015, show that using

¹⁵ <http://storm.apache.org/> ¹⁶ <http://spark.apache.org/> ¹⁷ <https://matrix.org/docs/spec/>

Bitcoin over the Tor network opens a new attack surface [5]. Solutions to the tracing and identification problems identified by these researchers have been proposed and implemented in alternative cryptocurrency protocols: mixing using ring signatures and zero-knowledge proofs [36, 51].

A proposed improvement to standard ring signatures is the CryptoNote protocol, which prevents tracing assets back to their original owners by mixing together incoming transactions and outgoing transactions. This service hides the connections between the prosumers and the addresses. Mixing requires the possibility to create new wallets at will and the existence of a sufficient number of participants in the network. Monero is an example of a cryptocurrency that provides built-in mixing services by implementing the CryptoNote protocol [39]. There are however alternative implementations of mixing protocols such as CoinShuffle [46] or Xim [6]. A variant of ring signatures, group signatures, were first introduced by Chaum and van Heyst, 1991, [9] and then built upon by Rivest *et al.*, 2001 [45]. The basis for anonymity in the CryptoNote protocol, however, is a slightly modified version of the *traceable ring signature* algorithm by Fukisaki and Suzuki, 2007 [17]. This allows a member of a group to send a transaction so that it is impossible for a receiver to know any more about the sender than that it came from a group member without the use of a central authority.

Some newer cryptocurrencies, such as Zerocoin [36], provide built-in mixing services, which are often based on cryptographic principles and proofs.

7.5.3 Smart Meters' Privacy. Most works discussing privacy look at it from the context of smart meters. McDaniel and McLaughlin discuss privacy concerns due to energy-usage profiling, which smart grids could potentially enable [34]. Efthymiou and Kalogridis describe a method for securely anonymizing frequent electrical metering data sent by a smart meter by using a third party escrow mechanism [12]. Tan *et al.* study privacy in a smart metering system from an information theoretic perspective in the presence of energy harvesting and storage units [49]. They show that energy harvesting provides increased privacy by diversifying the energy source, while a storage device can be used to increase both energy efficiency and privacy. However, transaction data from energy trading may provide more fine-grained information than smart meter based usage patterns [20].

8 CONCLUSIONS

In this paper, we described TRANSAX, a decentralized platform for implementing energy exchange mechanisms in a microgrid setting. Building on top of blockchains, we obtained decentralized trust and consensus capabilities, which prevent malicious actors from tampering with the shared system state. We demonstrated that the system is safe from anonymous offers, satisfying the seemingly conflicting goals of safety and privacy. Using our hybrid-solver approach, which combines a smart-contract based validator with an external optimizer, we showed that we can clear offers securely, efficiently, and resiliently. We also described a simulation testbed that allows us to test market platforms on a simulated microgrid, and we presented results showing how the platform can reduce the load on the bulk grid.

8.1 Future Work

In this work, we focused on satisfying the requirements of safety, efficiency, privacy, and resilience. We have not chosen a specific approach for *setting the clearing prices* for the prosumers' trades since the economics of setting the clearing prices is an orthogonal problem. Friedman and Rust [16] provide a survey of these mechanisms for governing trade, to which they refer as market institutions. One of the most commonly used mechanisms is the double auction. Note that we cannot apply the double auction directly because of the different time-interval attributes that the offers may specify. Prior work has extended the double auction to allow for multiple attributes; however, they typically (e.g. [3]) require a function to combine the attributes into a single value, which is then

used to order the offers. The difficulty of this approach is in identifying a meaningful function. A more straightforward approach is to perform the feasibility matching as we have presented, and then for each interval, use a double auction to set the clearing price for the matched offers. This approach provides a straightforward solution to the problem of setting clearing prices; however, it is not obvious whether it will preserve the properties that a simple double auction has, such as incentive compatibility. We leave the investigation of these mechanisms and how they are impacted by privacy to future work.

A second extension is to enable prosumers to *update or cancel offers*. The current formulation can support updating offers as long as the updates do not invalidate previous solutions; for example, a selling offer can increase the amount of energy to be sold or augment the set of intervals in which energy could be produced. To support restrictive changes or cancelling offers, we would need to introduce a deadline for when offers could no longer be updated or cancelled. Solvers could then wait for this deadline, and start working only after the deadline.

Finally, we have not yet discussed how the solvers are *incentivized to match the offers*. One option is to provide some reward for each solution that is finalized. Another is to expect the prosumers to submit solutions that are beneficial to themselves. These and other potential incentive mechanisms have yet to be evaluated.

REFERENCES

- [1] Merlinda Andoni, Valentin Robu, David Flynn, Simone Abram, Dale Geach, David Jenkins, Peter McCallum, and Andrew Peacock. 2019. Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renewable and Sustainable Energy Reviews* 100 (2019), 143–174.
- [2] M Apostolaki, A Zohar, and L Vanbever. 2017. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*. 375–392. <https://doi.org/10.1109/SP.2017.29>
- [3] Gaurav Baranwal and Deo Prakash Vidyarthi. [n.d.]. A fair multi-attribute combinatorial double auction model for resource allocation in cloud computing. 108 ([n. d.]), 60–76. <https://doi.org/10.1016/j.jss.2015.06.025>
- [4] Simon Barber, Xavier Boyen, Elaine Shi, and Ersin Uzun. 2012. *Bitter to Better — How to Make Bitcoin a Better Currency*. Springer Berlin Heidelberg, Berlin, Heidelberg, 399–414. https://doi.org/10.1007/978-3-642-32946-3_29
- [5] A Biryukov and I Pustogarov. 2015. Bitcoin over Tor isn’t a Good Idea. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*. 122–134. <https://doi.org/10.1109/SP.2015.15>
- [6] George Bissias, A Pinar Ozisik, Brian N Levine, and Marc Liberatore. 2014. Sybil-Resistant Mixing for Bitcoin. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society (WPES ’14)*. ACM, New York, NY, USA, 149–158.
- [7] Alain Brenzikofer, Arne Meeuw, Sandro Schopfer, Anselma Wörner, and Christian Dürr. 2019. Quartierstrom: A Decentralized Local P2P Energy Market Pilot On A Self-Governed Blockchain. In *Proceedings of the 25th International Conference on Electricity Distribution*. 5.
- [8] David P Chassin, K Schneider, and C Gerkenmeyer. 2008. GridLAB-D: An open-source power systems modeling and simulation environment. In *2008 IEEE/PES Transmission and Distribution Conference and Exposition*. IEEE, 1–5.
- [9] David Chaum and Eugène van Heyst. 1991. Group Signatures. In *Proceedings of the 10th Annual International Conference on Theory and Application of Cryptographic Techniques (EUROCRYPT)*. 257–265.
- [10] William Cox and Toby Considine. 2013. Structured energy: Microgrids and autonomous transactive operation. In *Proceedings of the 2013 IEEE PES Innovative Smart Grid Technologies (ISGT)*. 1–6.
- [11] O Dag and B Mirafzal. 2016. On stability of islanded low-inertia microgrids. In *2016 Clemson University Power Systems Conference (PSC)*. IEEE, Clemson, SC, USA, 1–7.
- [12] Costas Efthymiou and Georgios Kalogridis. 2010. Smart grid privacy via anonymization of smart metering data. In *Proceedings of the 1st IEEE International Conference on Smart Grid Communications (SmartGridComm)*. Gaithersburg, MD, USA, 238–243.
- [13] Scott Eisele, Abhishek Dubey, Gabor Karsai, and Srdjan Lukic. 2017. Transactive Energy Demo with {RIAPS} Platform. In *8th International Conference on Cyber Physical Systems (ICCP)*. IEEE, Pittsburgh, PA, USA, 91.
- [14] Scott Eisele, Aron Laszka, Anastasia Mavridou, and Abhishek Dubey. 2018. {SolidWorx}: A Resilient and Trustworthy Transactive Platform for Smart and Connected Communities. In *2018 IEEE International Conference on Blockchain (Blockchain-2018)*. IEEE, Halifax, NS, Canada, Canada, 1263–1272.
- [15] Scott Eisele, Istvan Madari, Abhishek Dubey, and Gabor Karsai. 2017. RIAPS: Resilient Information Architecture Platform for Decentralized Smart Systems. In *2017 IEEE 20th International Symposium on Real-Time Distributed*

- Computing (ISORC)*. IEEE, IEEE, Toronto, ON, Canada, 125–132.
- [16] Daniel Friedman. [n.d.]. *The Double Auction Market Institutions, Theories, and Evidence: Proceedings of the Workshop on Double Auction Markets held June, 1991 in Santa Fe, New Mexico* (1 ed.). Routledge. <https://doi.org/10.4324/9780429492532>
 - [17] Eiichiro Fujisaki and Koutarou Suzuki. 2007. *Traceable Ring Signature*. Springer Berlin Heidelberg, Berlin, Heidelberg, 181–200. https://doi.org/10.1007/978-3-540-71677-8_13
 - [18] J. Guerrero, A. C. Chapman, and G. Verbič. 2018. Decentralized P2P Energy Trading under Network Constraints in a Low-Voltage Network. *IEEE Transactions on Smart Grid* (2018), 1–1. <https://doi.org/10.1109/TSG.2018.2878445>
 - [19] Bing Huang, Alvaro A. Cardenas, and Ross Baldick. 2019. Not Everything is Dark and Gloomy: Power Grid Protections Against IoT Demand Attacks. In *28th USENIX Security Symposium*. USENIX Association, 1115–1132.
 - [20] A Hussain, V H Bui, and H M Kim. 2017. A Resilient and Privacy-Preserving Energy Management Strategy for Networked Microgrids. *IEEE Transactions on Smart Grid* PP, 3 (2017), 2127–2139.
 - [21] S Jebri, M Abid, and A Bouallegue. 2017. STAC-protocol: Secure and Trust Anonymous Communication protocol for IoT. In *2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC)*. 365–370.
 - [22] J. Kang, R. Yu, X. Huang, S. Maharjan, Y. Zhang, and E. Hossain. 2017. Enabling Localized Peer-to-Peer Electricity Trading Among Plug-in Hybrid Electric Vehicles Using Consortium Blockchains. *IEEE Transactions on Industrial Informatics* 13, 6 (Dec 2017), 3154–3164. <https://doi.org/10.1109/TII.2017.2709784>
 - [23] Amandeep Kaur, Jitender Kaushal, and Prasenjit Basak. 2016. A review on microgrid central controller. *Renewable and Sustainable Energy Reviews* 55 (2016), 338–345. <https://doi.org/10.1016/j.rser.2015.10.141>
 - [24] M. Khorasany, Y. Mishra, and G. Ledwich. 2017. Auction based energy trading in transactive energy market with active participation of prosumers and consumers. In *Proceedings of the 2017 Australasian Universities Power Engineering Conference (AUPEC)*. 1–6.
 - [25] Koen Kok and Steve Widergren. 2016. A society of devices: Integrating intelligent distributed resources with transactive energy. *IEEE Power and Energy Magazine* 14, 3 (2016), 34–45.
 - [26] Aron Laszka, Abhishek Dubey, Michael Walker, and Doug Schmidt. 2017. Providing privacy, safety and security in IoT-based transactive energy systems using distributed ledgers. In *Proceedings of the 7th International Conference on the Internet of Things (IoT '17)*. ACM, New York, NY, USA, 13:1–13:8. <https://doi.org/10.1145/3131542.3131562>
 - [27] A. Laszka, S. Eisele, A. Dubey, G. Karsai, and K. Kvaternik. 2018. TRANSAX: A Blockchain-Based Decentralized Forward-Trading Energy Exchanged for Transactive Microgrids. In *2018 IEEE 24th International Conference on Parallel and Distributed Systems (ICPADS)*. 918–927. <https://doi.org/10.1109/PADS.2018.8645001>
 - [28] Robert M. Lee, Michael J. Assante, and Tim Conway. 2016. *Analysis of the cyber attack on the Ukrainian power grid: Defense use case*. Technical Report. Electricity Information Sharing and Analysis Center (E-ISAC).
 - [29] Timo Lehtola and Ahmad Zahedi. 2019. Solar energy and wind power supply supported by storage technology: A review. *Sustainable Energy Technologies and Assessments* 35 (2019), 25–31.
 - [30] N W A Lidula and A D Rajapakse. 2011. Microgrids research: A review of experimental microgrids and test systems. *Renewable and Sustainable Energy Reviews* 15, 1 (2011), 186–202. <https://doi.org/10.1016/j.rser.2010.09.041>
 - [31] Peipeng Liu, Lihong Wang, Qingfeng Tan, Quangang Li, Xuebin Wang, and Jinqiao Shi. 2014. Empirical Measurement and Analysis of I2P Routers. *Journal of Networks* 9 (2014), 2269–2278.
 - [32] Power Ledger Pty Ltd. 2019. Power Ledger Whitepaper. <https://www.powerledger.io/wp-content/uploads/2019/05/power-ledger-whitepaper.pdf> (Accessed on 28 August 2019).
 - [33] Bodhisattwa P Majumder, M Nazif Faqiry, Sanjoy Das, and Anil Pahwa. 2014. An efficient iterative double auction for energy trading in microgrids. In *2014 IEEE Symposium on Computational Intelligence Applications in Smart Grid (CIASG)*. IEEE, Orlando, FL, USA, 1–7. <https://doi.org/10.1109/CIASG.2014.7011556>
 - [34] Patrick McDaniel and Stephen McLaughlin. 2009. Security and privacy challenges in the smart grid. *IEEE Security & Privacy* 7, 3 (2009), 75–77.
 - [35] Ronald B Melton. 2013. *Gridwise transactive energy framework*. Technical Report. Pacific Northwest National Laboratory, Richland, WA.
 - [36] I Miers, C Garman, M Green, and A D Rubin. 2013. Zerocoin: Anonymous Distributed E-Cash from Bitcoin. In *2013 IEEE Symposium on Security and Privacy*. IEEE, Berkeley, CA, USA, 397–411. <https://doi.org/10.1109/SP.2013.34>
 - [37] S J Murdoch and G Danezis. 2005. Low-cost traffic analysis of Tor. In *2005 IEEE Symposium on Security and Privacy*. 183–195. <https://doi.org/10.1109/SP.2005.12>
 - [38] Leonardo Neumeyer, Bruce Robbins, Anish Nair, and Anand Kesari. 2010. S4: Distributed stream computing platform. In *Data Mining Workshops (ICDMW), 2010 IEEE International Conference on*. IEEE, 170–177.
 - [39] Shen Noether. 2015. Ring Signature Confidential Transactions for Monero. Cryptology ePrint Archive, Report 2015/1098.
 - [40] HVS Kumar Nunna and Suryanarayana Doolla. 2013. Multiagent-Based Distributed-Energy-Resource Management for Intelligent Microgrids. *IEEE Transactions on Industrial Electronics* 60, 4 (April 2013), 1678–1687.

- [41] Farrokh A Rahimi and Ali Ipakchi. 2012. Transactive energy techniques: closing the gap between wholesale and retail markets. *The Electricity Journal* 25, 8 (2012), 29–35.
- [42] Tom Randall. 2015. The way humans get electricity is about to change forever. Bloomberg.
- [43] Michael G Reed, Paul F Syverson, and David M Goldschlag. 1998. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communications* 16, 4 (1998), 482–494.
- [44] Fergal Reid and Martin Harrigan. 2013. An Analysis of Anonymity in the Bitcoin System. In *Security and Privacy in Social Networks*. 197–223.
- [45] Ronald L Rivest, Adi Shamir, and Yael Tauman. 2001. How to Leak a Secret. In *Proceedings of the 7th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*. 552–565.
- [46] Tim Ruffing, Pedro Moreno-Sanchez, and Aniket Kate. 2014. CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin. In *Proceedings of the 19th European Symposium on Research in Computer Security (ESORICS)*. 345–364.
- [47] Douglas C Schmidt, Jonathan White, and Christopher D Gill. 2014. Elastic Infrastructure to Support Computing Clouds for Large-scale Cyber-Physical Systems. In *2014 IEEE 17th International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC)*. 56–63.
- [48] J W Simpson-Porco, Q Shafiee, F Dörfler, J C Vasquez, J M Guerrero, and F Bullo. 2015. Secondary Frequency and Voltage Control of Islanded Microgrids via Distributed Averaging. *IEEE Transactions on Industrial Electronics* 62, 11 (November 2015), 7025–7038.
- [49] Onur Tan, Deniz Gunduz, and H Vincent Poor. 2013. Increasing smart meter privacy through energy harvesting and storage devices. *IEEE Journal on Selected Areas in Communications* 31, 7 (2013), 1331–1341.
- [50] Taha Selim Ustun, Cagil Ozansoy, and Aladin Zayegh. 2011. Recent developments in microgrids and example cases around the world – A review. *Renewable and Sustainable Energy Reviews* 15, 8 (2011), 4030–4041.
- [51] Nicholas van Saberhagen. 2012. *CryptoNote v 2.0*. Technical Report. 3 pages. [https://cryptonote.org/whitepaper\[_\]v2.pdf](https://cryptonote.org/whitepaper[_]v2.pdf)
- [52] Peter Volgyesi, Abhishek Dubey, Timothy Krentz, Istvan Madari, Mary Metelko, and Gabor Karsai. 2017. Time Synchronization Services for Low-cost Fog Computing Applications. In *28th International Symposium on Rapid System Prototyping (RSP)*. IEEE, New York, NY, USA, 57–63.
- [53] Jian Wang, Qianggang Wang, Niancheng Zhou, and Yuan Chi. 2017. A Novel Electricity Transaction Mode of Microgrids Based on Blockchain and Continuous Double Auction. *Energies* 10, 12 (2017). <https://doi.org/10.3390/en10121971>
- [54] Gavin Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper* 151 (2014).
- [55] Anselma Wörner, Arne Meeuw, Liliane Ableitner, Felix Wortmann, Sandro Schopfer, and Verena Tiefenbeck. [n.d.]. Trading solar energy within the neighborhood: field implementation of a blockchain-based electricity market. 2 ([n. d.]). Issue S1. <https://doi.org/10.1186/s42162-019-0092-0>
- [56] M Yazdanian and A Mehrizi-Sani. 2014. Distributed Control Techniques in Microgrids. *IEEE Transactions on Smart Grid* 5, 6 (November 2014), 2901–2909. <https://doi.org/10.1109/TSG.2014.2337838>
- [57] F Yucel, E Bulut, and K Akkaya. 2018. Privacy Preserving Distributed Stable Matching of Electric Vehicles and Charge Suppliers. In *Proceedings of the 2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*. 1–6.
- [58] Kim Zetter. 2016. Inside the Cunning, Unprecedented Hack of Ukraine’s Power Grid. WIRED, <https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/>.
- [59] Chenghua Zhang, Jianzhong Wu, Yue Zhou, Meng Cheng, and Chao Long. 2018. Peer-to-Peer energy trading in a Microgrid. *Applied Energy* 220 (2018), 1–12. <https://doi.org/10.1016/j.apenergy.2018.03.010>
- [60] Liang Zhao, Kwangho Park, and Ying-Cheng Lai. 2004. Attack vulnerability of scale-free networks due to cascading breakdown. *Physical Review E* 70, 3 (sep 2004), 35101. <https://doi.org/10.1103/PhysRevE.70.035101>
- [61] Shengnan Zhao, Beibei Wang, Yachao Li, and Yang Li. 2018. Integrated Energy Transaction Mechanisms Based on Blockchain Technology. *Energies* 11, 9 (2018). <https://doi.org/10.3390/en11092412>