A Block-Free Distributed Ledger for P2P Energy Trading: Case with IOTA?*

Joon Park[1], Ruzanna Chitchyan[1], Anastasia Angelopoulou[2], and Jordan Murkin[1]

[1] University of Bristol, BS8 1UB, UK
{jp17807,r.chitchyan,jordan.murkin}@bristol.ac.uk
[2] Columbus State University, Columbus, GA, USA
angelopoulou_anastasia@columbusstate.edu

Abstract. Across the world, the organisation and operation of the electricity markets is quickly changing, moving towards decentralised, distributed, renewables-based generation with real-time data exchange-based solutions. In order to support this change, blockchain-based distributed ledgers have been proposed for implementation of peer-to-peer energy trading platform. However, blockchain solutions suffer from scalability problems as well as from delays in transaction confirmation. This paper explores the feasibility of using IOTA's DAG-based block-free distributed ledger for implementation of energy trading platforms. Our agent-based simulation research demonstrates that an IOTA-like DAG-based solution could overcome the constraints that blockchains face in the energy market. However, to be usable for peer-to-peer energy trading, even DAG-based platforms need to consider specificities of energy trading markets (such as structured trading periods and assured confirmation of transactions for every completed period).

Keywords: blockchain · peer to peer energy trading platform · DAG-based distributed ledger · block-free ledger · IOTA · agent-based simulation.

1 Introduction

In the current energy market, utility companies act as intermediaries between householders and the market, purchasing any excess generation that households produce. This is shown in Fig. 1.a. In contrast to this, a **peer-to-peer (p2p) energy market** enables any two individuals to directly buy from and sell to each other, without the utility-intermediaries [1], as shown in Fig. 1.b. Such households can be both prosumers (i.e., producing and consuming own electricity, as well as selling the excess to others), or only consumers (if they don't own any generation facilities). The key advantages here are in providing avenues for:

* This work is supported by the UK EPSRC funding for Refactoring Energy Systems (EP/R007373/1) and Household Supplier Energy Market (EP/P031838/1) projects.

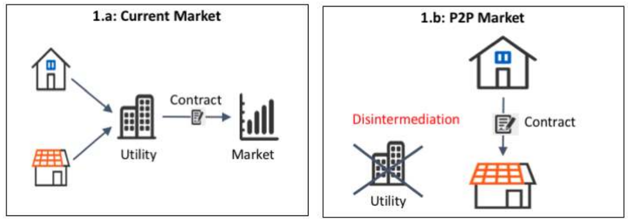

Fig. 1. Energy Market Disintermediation(from [2])

- Additional income streams to households with small-scale generation - via sale of their excess energy to other peers in the market[3]. Moreover, the price of the locally generated energy is likely to be more competitive than that of the grid supply, as the microgenerators would have lower generation costs, and no intermediation fees paid to utilities[4].
- Additional (non-monetary) value proposition of small-scale generation and energy storage - the microgenerators not only get return on their investment into generation equipment, but also support the energy needs of local communities, contributing to efforts on decarbonisation and energy security.
- Increased control over source/destination of supply - consumers are able to express their preferences on energy purchase: do they wish to buy solar or wind, from the local producer or from the cheapest supplier; do producers wish to donate their excess generation to the local school or sell it to the highest bidder? All these options become viable when peers directly buy and sell from/to each other.

Such an energy system, however, requires **a digital platform** which will remove intermediation from the utilities (so that control is fully retained by the market participants), advertise the sale and purchase orders between the trading parties, undertake matching of these orders, based on the users' preferences, ensure security of the transactions, transparency of the trades, and accountability of the transaction participants.

Recently a number of researchers have advocated use of the blockchain ledgers to create such a p2p energy trading platform.This is due to the functionality that the blockchains enable. They:

- provide full support for **distributed, decentralised data storage and processing**. Thus, there is no need to use and pay for any centralised data storage organisations and facilities.

[3] Currently in the UK the excess generation must be sold back to the utility provider at a set rate (the so called feed in tariff set by the UK government). However, for many types of household generation this scheme will cease as of March 31 2019.

[4] Please note: each locality remains interconnected with the grid, the energy costs will still include grid connection and maintenance charges. This is because the households wish to be insured against the intermittency of the renewables-based generation, and the grid provides such insurance and balancing services.

- **remove the need for trusting a third party intermediary**, as the blockchain ledgers rely on agreement of the majority of the participants on the state of the chain, not on any external party;
- maintain tampering-resistant and **accountable records** of transactions.

Yet, the unpermissioned blockchains [27] (i.e., ledgers where all participants have unrestricted right to participate in transaction validation and data access - which is critical for an open and trusted p2p trading environment) also suffer from a number of **drawbacks**, such as:

- Scalability limitations when handling a large number of transactions, due to block size constraints and increasingly high transaction processing fees,
- Latency of transaction confirmation (particularly for low-fee transactions) due to low incentives to the miners to include such transactions into blocks.

In this paper, employing the IOTA ledger as a sample, we explore the feasibility of using the Directed Acyclic Graph-based block-free ledgers (abbreviated to DAG-BF) for p2p energy trading. Unlike blockchains, DAG-BF ledgers do not structure transaction records into blocks, and do not rely on specialist miners to process the transactions. Thus, we expect that the above mentioned blockchain limitations would be addressed. Moreover, the DAG-FB ledgers maintain transparent and decentralised transaction records, which fit with the needs of the p2p energy trading domain.

This paper proceeds by outlining the current state of the research and developments of blockchain solutions in the energy sector (section 2). An introduction to the key characteristics of the DAG-BF ledgers is also presented (section 2). Section 3 of the paper presents the p2p energy trading model and outlines the study design. Employing agent-based simulation, this study then sets out to investigate the feasibility of using DAG-BF ledgers for p2p energy trading. The findings of this study suggest that such ledgers could indeed be used for p2p-based energy trading (section 4). A discussion on how the peculiarities of the energy market would align with the present findings is also presented (section 4).

2 Distributed Ledgers in the Energy Sector: a Background Overview

A distributed ledger is a database architecture which facilitates peer-to-peer transactions in a distributed and decentralised way without the need for an intermediary or a centralised authority [27]. The following subsections provide a brief overview of blockchains and DAG-FB distributed ledgers in the energy sector.

2.1 Blockchains in the Energy Sector

Blockchain is a distributed ledger which records transactions, agreements, contracts, and sales [27, 7]. Here, a set of transactions is collected into a *block*. Each block of transactions is then validated by specialised peers on the network, called *miners*. Miners are rewarded for the validation effort with *transaction fees*. The validated transactions are recorded into the *ledger* (or chain).

The idea of using blockchains in the energy sector is becoming increasingly popular, as shown by the growing number of pilots and research projects [8]. It is often considered to be a game-changer for the energy industry [3, 27], as it has the potential to enable transition to low-carbon sustainable energy systems [4]; foster innovation in development of IoT platforms [5], digital applications for P2P energy trading and smart grids [8, 9].

Various electricity and gas distributors in different countries (such as Vector in New Zealand, Vattenfall in Sweden, EDF Energy and Verv in the UKto name a few) are already testing blockchain platforms for local p2p energy markets [8].

The research community has also explored the use of blockchain ledgers in P2P energy trading. Mengelkamp at el. [10] simulated a local energy market of 100 residential households where consumers and prosumers can trade energy within their community on a private blockchain platform. Murkin et al. [11] proposed a p2p electricity trading platform under a blockchain scheme to automatically buy and sell electricity in each household as microgeneration increases. Pop at el. [12] used a blockchain mechanism to manage the demand response in smart grids. Oh et al. [13] implemented an energy-trading system using MultiChain and demonstrated that transactions worked correctly over blockchain. The use of blockchains for sharing of renewable electricity through smart contracts was studied in NRGX-Change [22] and the Crypto-Trading projects [23].

However, the structure of blockchain-based solutions has recently been critisised due to the difference between mining and other nodes, as well as block size restrictions [10, 13]. The miners (i.e., block validator nodes) are motivated by transaction processing fees for including a transaction into the block-to-be-validated. Consequently, the transactions willing to pay higher fees are given priority for inclusion into the blocks, pushing the fees to increase as the number of waiting transactions increases. Transactions with low allocated processing fees may remain in the queue of unconfirmed transactions for a long time, as the block sizes are limited and higher-fee transactions are always chosen first.

2.2 DAG-based Block Free Distributed Ledgers

In a block-free ledger the individual transactions are directly introduced into the ledger (without aggregation into blocks). The newly introduced transactions also cross-verify other transactions, thus carrying out the task which was done by the dedicated miners in the blockchain. Thus, the block-free distributed ledger removes the distinction between miner and participant nodes. Here all nodes of the network must participate in the transaction approval.

In many current block-free ledgers (BF) [16–18] the cross-validating transactions are structured into a Directed Acyclic Graph (DAG). Such a DAG consists of vertices and edges, where each vertex represents a transaction and each directed edge a reference. The referencing edges validate and approve the transactions to which they point. The DAG serves as a truly distributed ledger, which reaches consensus by accumulating information about the state of the network.

It must be noted that a DAG structure can also be used within blockchains, for instance, Ethereum [19] employs a DAG where *blocks of transactions* comprise the vertices. The problems of block size, transaction validation latency, and use of dedicated miners, however, remain. Thus, we focus on DAG-based block-free (DAG-BF) alternative in our study.

Several DAG-BF cryptocurrencies have recently gained recognition, such as RaiBlocks [16], Byteball [17] and IOTA [18]. These differ from each other in the details of implementation and consensus protocols. For instance, IOTA requires that each transaction is referenced by two other transactions for verification, while, for Byteball, references to a number of trusted nodes are necessary. IOTA achieves consensus via the cumulative proof-of-work of confirmed transactions, while, in RaiBlocks, consensus is achieved via balance-weighted vote on conflicting transactions. Yet, they all have a common set of characteristics which are relevant for implementing a distributed ledger-based platform for a p2p energy trading market:

1. transactions are processed *individually*, without block formation, which overcomes the processing latency due to block size constraint;
2. processing is carried out *asynchronously*, "upon arrival" of each new transaction, tackling the delay of block formation;
3. each network participant is also a validator, *without distinction between mining and other nodes*;
4. newly arriving transactions are added as leaf nodes into a DAG, and are to be *confirmed by accumulating references* within the DAG.

In this work IOTA is used as a sample DAG-BF ledger to investigate the feasibility of using such ledgers in p2p energy trading. While the details of the simulation are, by necessity, aligned with the IOTA specifics (e.g., referencing 2 parent nodes, cumulative weight calculation method), the results that relate to the above DAG-BF ledger characteristics could be considered of wider relevance.

2.3 IOTA

IOTA is a DAG-BF distributed ledger; its DAG is called the Tangle [20]. When a new transaction enters the Tangle, it selects two existing transactions to approve, and an edge is created between the new transaction and each of its selected predecessors. The new transaction then approves the two selected transactions (by solving a cryptography puzzle that links it to its approved transactions) and waits for another transaction to approve it. (e.g., see Fig. 2).

An unconfirmed transaction in a DAG is called a tip (e.g., the transactions A, B, X, C, E in Fig. 2, as these have less than 2 incoming transactions confirming them). The average time that a transaction remains unconfirmed in the Tangle (i.e., *transaction confirmation latency*) and the *number of unconfirmed transactions* (i.e. tips) at any given time are key parameters when considering the use of Tangle as a candidate for a p2p energy trading platform.

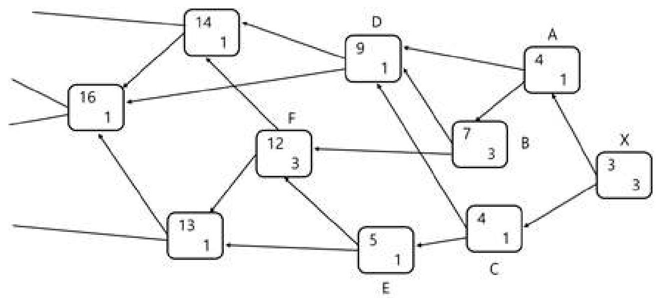

Fig. 2. View of IOTA Tangle (from [20]. Boxes represent transactions, number at the bottom of a box is own weight, number at the top is the cumulative weight.)

Both the *transaction confirmation latency* and the *number of unconfirmed tips at any given time* are heavily dependent on how new transactions select the tips to validate, i.e., on their tip-selection algorithm [20]. Additionally the *number of unconfirmed tips* is also dependent on the rate of arrival of new transactions into the Tangle. Examples of tip selection algorithms are:

– *Uniform random selection*, whereby each new transaction randomly selects two tips to confirm, (without having to traverse the graph);
– *Unweighted random walk*, whereby starting from the genesis node, the walker chooses which transaction to move to with equal probability; and
– *Weighted random walk* where each transaction is assigned a weight (e.g., 1 for D in Fig. 2) and a cumulative weight [20] (e.g., 9 for D in Fig. 2). The weight defines how much work has been invested into each transaction by the issuing node. Cumulative weight is the sum of the weights of the given transaction and all other transactions that directly or indirectly reference the given one. For instance, cumulative weight for D in Fig 2 is 9, and it comprises the weight of D (1), weights of B(3), A (1), X (3) which indirectly references D through A, and C (1). Under the weighted random walk selection algorithm, a validating node selects a tip based on the tip's weight.

The complexity of the weights is introduced into the tip selection algorithms so as to penalise the so-called lazy (and/or malicious) nodes [20]. Lazy nodes send new transactions and attach them to already approved transactions, instead of approving new ones. Such behaviour saves computation time for the lazy nodes, but results in a larger number of unconfirmed transactions. Neither uniform random selection nor unweighted random walk algorithms can discriminate against the lazy nodes. The weighted random walk, on the other hand, can use the cumulative weights as an indicator for selection of 'honest' transactions. IOTA does this through use of the α parameter that can be set between 0 and 1, biasing the graph traversing towards the selection of higher-weighted transactions [20]. Thus (as shown in [20]) if y approves x then the transition probability

P_{xy} is:

$$P_{xy} = \exp\left(-\alpha(H_x - H_y)\right) / \sum_{z:z \to x} \exp\left(-\alpha(H_x - H_z)\right) \quad (1)$$

where H_x and H_y are the cumulative weights of nodes x and y respectively. If z approves x, then $z \to x$ and $\alpha > 0$ needs to be chosen. When α is set to 0 the algorithm reverts to unweighted random walk, as then there is no bias towards choosing the path through transactions with a higher cumulative weight. As α approaches 1, the walker will always choose only the path through the transaction with the highest cumulative weight, continuously increasing the weight of transactions in that path (by adding new indirect references).

3 Modelling a Block-Free P2P Energy Trading Platform

In this exploratory study we wish to investigate whether p2p energy trading would be feasible on a DAG-BF ledger with respect to transaction confirmation latency and large volume of transactions to be processed. As noted before, we opt to use IOTA as the sample DAG-BF ledger, as it is one of the most widely publicised such ledgers. Yet IOTA's current implementation is also strongly criticised [28, 30, 29]. For instance, presently IOTA relies on a centralised coordinator[5] node, which renders it a centrally controlled network [30] (though there are claims that the coordinator will be removed imminently [31]). It could also suffer from replay vulnerability [29].

In order to abstract from the specifics of the current implementation, while retaining the key characteristics of IOTA's DAG-FB ledger, we turn to an agent-based simulation (instead of on-chain implementation), assuming that the coordinator-free version of the ledger is in place. Agent-based (AB) modelling is particularly suitable for the present research as it allows us to focus on the defining properties of DAG-BF.

Thus, we set the characteristics and behavioural rules for each individual undertaking the transactions (i.e., agent) in the simulation, then observe the collective impact of these behaviours and interactions among agents [22, 23] and their impact on the ledger. To build an AB model, we must detail its constituent parts, i.e.:

- **Agents**, who are defined as heterogeneous entities with different characteristics and individual behaviours. They are situated in an environment and perform actions.
- **Interactions** between agents change the agents' state;
- **Environment** is the space in which agents are located (e.g., longitude and latitude), and the rules under which they operate (e.g., the excess energy that has been generated is to be sold).

[5] Note: in case of the energy sector such a coordinator may be acceptable, if run by the energy regulator providing governance to the system.

3.1 Study Design

To develop a P2P energy trading platform, we must first consider the characteristics of the energy market (i.e., the environment within which the actors operate), including trade organisation, and the properties of traders.

Market Structure: Trade Periods. On the current energy market, energy is bought and sold per 30 min. intervals. This structure is dictated by the nature of the market itself, as the amounts of electricity generation and consumption (and so its prices) vary due to:

- *external environmental conditions* (especially for renewables-based generation) as, for instance, PV panels generate more in the summer than in the winter and consumers use more electricity in cold weather, as well as
- *time of day*, e.g., PV panels have highest output around midday; while most individuals wake up and have breakfast between 6.30 AM and 9 AM on weekdays, (causing increased demand in the morning), yet
- the electricity *grid must be balanced for every time period*, which requires for the peak-time energy use (when grid is under stress to meet the high consumption requirements) to be more expensive than off-peak.

Thus, we too keep to the period-based trading, where the trading is carried out every t period (where t can be 30 min, or less). During each current period the sellers and buyers publish their desired sell/buy request for the next period. At the end of the current period the buy and sell requests are matched and recorded as transactions into the ledger. A new period then begins. For simplicity we do not consider cases where advertised sell/buy requests are not satisfied and further settlement is required, as this does not affect the ledger's scalability or latency, but depends on the trading and settlement algorithm used [24]. Similarly, though the matching algorithm allows for partial trading, where one individual buys/sells to many in a given trading period, to fill his/her order, these are simply extra transactions to the ledger, and are not further discussed. The key result of such trade structuring decision is that all the sell/buy requests are processed together, and *their results are released for committing into the ledger at the same time.*

Households and Interactions. We model individual households as *automated software buyer or seller agents* that express their trading preferences, and can generate (as producers) or use (as consumers) energy. The households periodically (e.g., every 30 min) advertise the amount of generation (sell requests) and use (buy requests) they provide/require. A p2p trading algorithm then calculates a stable match for each of the buyer-seller pair of agents, while taking into account their individual preferences (e.g., as much as possible, buy only solar energy, or sell to local buyers). The national grid acts as the default seller/buyer where the p2p market under/over produces. Once the stable match is found (i.e., there is no buyer/seller pair that would prefer to be matched to a different partner), transactions are sent to the distributed ledger by the buyers. The attributes of the agents are illustrated in Table 1.

Table 1. Agent Structure.

Properties	Definition
Agent ID	Unique agent identifier
Agent location	Agent location in terms of latitude and longitude
Amount of electricity	Amount of electricity to buy or sell for every trading period
Generation Type	Generation type (i.e., solar PV, wind, anaerobic digestion, hydro, and micro CHP)
Distance preference	Distance preference for the trade
Price preference	Minimum price (willing to sell) for seller agent or maximum price (willing to buy) for buyer agent

A seller agent's properties are the amount of electricity to sell, the location, the generation type, the distance preference (how far the buyer can be located), and the minimum price at which the seller is willing to sell. A buyer agent's properties are the amount of electricity to buy, his/her location, distance preferences, the maximum price at which the buyer is willing to buy, and the preferred generation-type to purchase (as a priority list, ranked from 1 to 5).

The agents' attributes are used to express agent preferences and to influence with which other agents a given agent will trade. Thus, the distance preferences of each seller/ buyer are used to determine which potential trading partners are located within his/her preferred distance. If a party is located outside the distance preference, then the matching score to this party will be lowered. The seller's minimum price at which (s)he is willing to sell is the minimum sales price per 1 kWh of electricity, and the buyer's maximum price at which (s)he is willing to buy should be greater than the seller's minimum sales price in order for these two agents to enter into a transaction.

The p2p electricity trading process starts with sellers and buyers publishing their desired sell/buy requests. Then, the trading algorithm[6] scores each buyer-seller combination, and ranks the matches based on their scores. The highest scoring pairs are matched. The buyer then creates a transaction, which is recorded into the ledger. The transaction is initiated by the buyer, as it is fully dependent on his/her willingness to pay. Each transaction contains the transaction ID, the buyer's ID, the seller's ID, the amount of electricity exchanged, the unit price, and the timestamp. This information about each transaction is saved in a file and is used to check the shape of the transaction graph using the GraphViz [26] visualisation tool. The AB model was developed with the AnyLogic [21] simulation tool[7].

[6] The details of the matching algorithm are not the focus of this study. The base algorithm is given in [24], though this study uses an extended version.

Model Setup Before the model can be executed, it needs to be set up[7]. The accounts are randomly generated for the simulation. We assume that buyers and seller of the p2p electricity market area are located in an arbitrary area in the UK, with latitude randomly chosen between 50.956870 and 52.438562, and longitude between -2.386779 and 0.292914. Sellers' generation types are uniformly distributed, and sellers generate from 5 to 10 kWh every trading period. Five types of generation are used: solar, wind, hydro, anaerobic digestion, and micro CHP. Distance preference is set for all accounts between 5-10km. For buyers, the maximum price to buy is set randomly between 14p and 16p per kWh, and the demand is set randomly between 1 kWh and 6 kWh. Buyers' generation-type preferences are set randomly. During model testing, a percentage of these accounts were assigned to be sellers. Simulations were run for 16 replications for each model, where the percentage of sellers in the market ranged from 5% to 20% in 5% increments, and the number of participants ranged from 500 to 3000 in increments of 500.

4 Findings and Discussion

After setting the number of sellers and buyers that participate in the electricity trading market, as well as their location, sales volume, and demand quantity, we can observe the changes in the metrics and monitor the transactions that occur among them. The metrics of this feasibility study are the transaction confirmation latency and the number of unconfirmed tips per trading period.

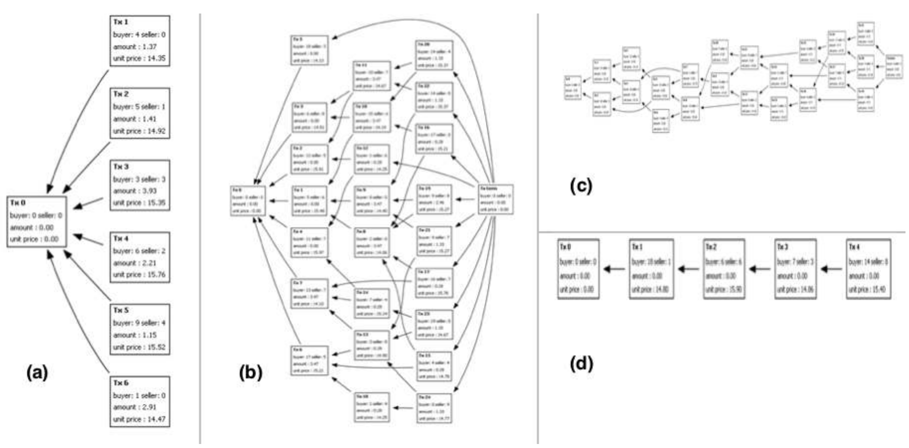

Fig. 3. Shapes of Tangle under various transaction release rates (the last nodes in parts b and c are termination nodes, see section 4.1).

---

[7] This model can be accessed via: https://cloud.anylogic.com/model/966f6846-62e0-460e-bf69-2a1b00317128?mode=SETTINGS from the Anylogic Cloud.

4.1 Impact of Market Structure

Trade Periods: As previously discussed, the nature of the energy market requires that trades are structured into fixed time intervals[8]. As the trades are agreed at the end of the trading period, the buyers' agents will simultaneously[9] send the transactions into the ledger. Simultaneous arrival of large groups of transactions into the ledger could result in several transactions simultaneously selecting and confirming the same tip with some other tips remaining unconfirmed for a longer time [20]. For instance, as shown in Fig. 3.a, all transactions released in the first group confirm solely the genesis node, while they themselves remain unconfirmed until transactions from the next group are released.

Fig. 3 illustrates the shape of the block-free distributed ledger under various (uniformly distributed) delay ranges in sending transactions (Fig. 3.(b) delay of 0.1s to 0.3s, (c) from 0.7s to 0.9s, (d) from 1.3s to 1.5s.) [10]. As the agents' delay in setting out transactions increases, the shape of the Tangle of the block-free distributed ledger converges to the shape of a chain.

Termination Nodes: The interval-based market structure in energy trading (e.g., per 30 min) implies that transactions which were not confirmed at the end of one trading period cannot be settled until the next trading period releases transactions. Yet, energy generation and consumption cannot be postponed until the next period during which transactions would be confirmed. To address this concern, we suggest the need for a termination node that would confirm all unconfirmed transactions at the end of each trading period. The work for creation of a termination node could be allocated to all network participants, with one or two participants randomly selected for such node generation at the end of each trading period. Two nodes will ensure that each unconfirmed transaction has two validators (as required per Iota's protocol). The termination nodes can then serve as the starting nodes for the DAG of the next set of trade transactions, as illustrated in Fig. 4. This structure both ensures the transactions are confirmed for each trading period and the DAG nodes are clearly allocated to each trading period. IOTA's current solution whereby all nodes maintain statistics on new transactions received from the neighbours [20] could still be sufficient in excluding transactions from "too lazy" nodes, and prevent a behaviour where nodes expect that the termination node "will confirm all transaction anyhow".

[8] While this study considers buy/sell requests published at period (t) for period (t+1), further periods (e.g., t+10, as different markets) can also be studied with this model.

[9] Specific implementations of the trading algorithms could vary the transaction release rate. For instance, when the ranked order algorithm is applied in trade matching, the matching is carried out in several cycles, and trades are agreed and released into the ledger in groups [24]. In this case, the following discussion relates to a single group release.

[10] As we focus on the shape of the ledger not transaction content, figure readability is not strictly necessary. Yet, if need be, the figures can be accessed via: https://jmp.sh/WNkIbZp

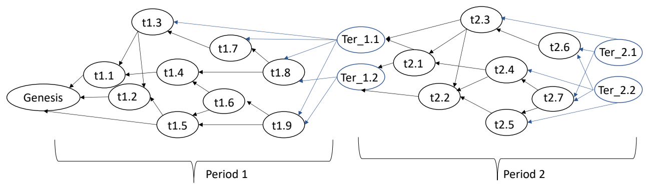

Fig. 4. Connecting two Directed Acyclic Graphs

4.2 Impact of Tip Selection Algorithm

As previously discussed, the transaction confirmation time and the number of unconfirmed tips present in the ledger at any given time are dependant on the tip selection algorithm [20].

To analyse the impact of the tip selection algorithms, the market in the AB model was simulated for 2,500 buyers and 375 sellers, with an average of 1,000 transactions generated and recorded into the block-free distributed ledger during the trading period. The delay of releasing the transactions into the ledger is uniformly distributed between 0.3 to 0.9 seconds.

Transaction confirmation time: The transaction confirmation time is defined as the difference between the time when a transaction is connected to the block-free distributed ledger and the time of its first transmission into the network.

The transaction confirmation times aggregated from the above discussed simulation are depicted in Fig. 5 when (a) uniform random selection, (b) unweighted random walk, and (c) weighted random walk are used.

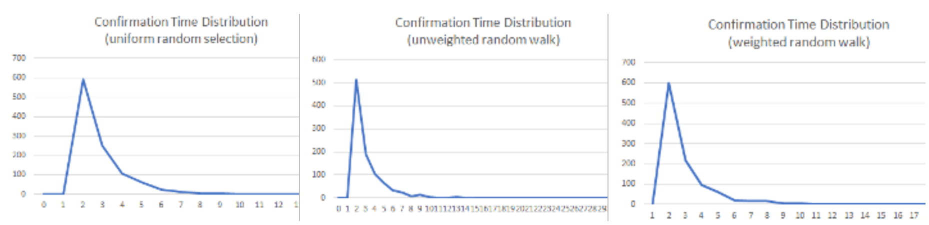

Fig. 5. Confirmation time distribution using (a) uniform random selection (b) unweighted random walk (b) weighted random walk.

While all 3 tip selection algorithms confirm between 500 and 600 transactions within 2 seconds, the average confirmation time varies from 2.26 for uniform random walk to 2.63 for unweighted and 7.91 sec for weighted random walk. The length of the last case is due to delays in confirming some outlier transactions.

Unconfirmed Transactions (Tips): The number of unconfirmed transactions recorded from the above noted simulation are depicted in Fig. 6 with: (a) uniform random selection, (b) unweighted random walk, and (c) weighted random walk. We observe that both uniform random selection and unweighted random walk

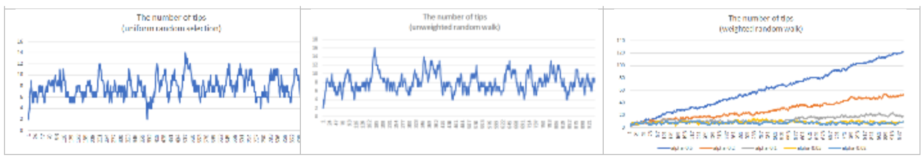

Fig. 6. Number of unconfirmed tips under (a) uniform random selection, (b) unweighted random walk, and (c) weighted random walk.

maintain stable ranges of unconfirmed tips (2 to 14 for random selection and 2 to 16 for unweighted random walk). However, under the weighted random walk, the number of unconfirmed tips could become divergent. Here the range stability depends on the alpha value (which was discussed in section 2.2 equation (1).

To analyse this, we recorded the number of unconfirmed transactions per unit time for the following α values: 0.02, 0.05, 0.1, 0.2 and 0.5(see Fig. 6.c). We noticed that as the value of alpha increases, the number of unconfirmed transactions shows an upwards trend, rather than being stably maintained. Thus, as alpha approaches 1, predominantly the transactions with larger cumulative weight are selected for confirmation, while other tips remain unconfirmed (see Fig. 6.c). On the other hand, when alpha approaches zero, the number of unconfirmed tips becomes stabilised.

Thus, though the weighted random walk tip selection algorithm can penalise lazy and malicious behaviour by isolating lazy nodes, our experiment suggests that it could also lead to increased transaction confirmation delay. As discussed in section 2.3, as α increases, the algorithm will converge to a single traversal path. Thus, any transactions that had chosen to confirm any tips outside of this "main path" will forever remain unconfirmed. Consequently, should α be wrongly chosen, the ledger will be destabilised.

5 Conclusions

Using the IOTA as an example, this paper explores the feasibility of utilising (IOTA-like) DAG-based block-free distributed ledgers for implementation of peer-to-peer energy trading platforms. This effort is motivated by the promise of the DAG-BF ledgers to remove the need for specialist miners/validators and their respective fees, threat of over-centralisation due to dominance of large mining pools, and the risk of long transaction confirmation delays for some (low value) transactions.

Our agent-based simulation experiments for a p2p energy market suggest that the functioning of this simulated market can indeed be successfully supported by a IOTA and similar DAG-FB ledgers. However, we also noted the need to carefully design the shape of the block-free ledger:

1. The peculiarity of the p2p energy market structure (i.e., the need to trade over discrete time periods, where all trade transactions agreed for one period could potentially be released simultaneously) necessitates a uniformly randomised process of releasing transactions into the ledger for each trading period.
2. As energy generation and consumption have to be continuously balanced in the grid, all trades for a given period should be confirmed and completed before the start of the next period. We have suggested to use a dedicated kind of nodes, (so called termination nodes) that guarantee confirmation of all 'honest' transactions at the end of each trading period. They help to both finalise sales for each period and to clearly structure records per each trading period. We also note that the termination nodes provide an ideal location for the so-called ledger maintenance tasks (such as pruning and confirmation) [20]. This, however, also leaves an open question as to if and how exactly should the termination nodes avoid confirmation of transactions from lazy nodes. For instance, one could explore the impact of choosing the lazy nodes as the main workers in the generation of the termination nodes, thus forcing them into active participation.
3. There is no single "best" tip selection algorithm, as some of the most frequently used solutions (i.e., uniform random walk and unweighted random walk) do not safeguard against potential lazy and malicious behaviour of some nodes at the expense of the others. The choice of the weighted random walk, on the other hand, could lead to an increased average transaction confirmation delay (due to some outlier tips and/or inappropriately set α).

Finally, we must also note that, though this simulation study suggests that IOTA (and similar DAG-FB ledgers) appear(s) to be feasible for implementation of p2p energy trading platforms, this conclusion cannot be fully verified without an actual implementation of such a platform. Such an implementation is our immediate future work.

References

1. Zhang, C., Wu, J., Long, C., Cheng, M.: Review of existing peer-to-peer energy trading projects. Energy Procedia (105), 2563–2568 (2017)
2. Murkin, J., Chitchyan, R., Ferguson, D. Towards peer-to-peer electricity trading in the UK. Presented at All Energy (2018) https://reedexpo.app.box.com/s/ plwhcfaqp6pnhxc8mcjznh7jtkevg9h1/file/292636529562
3. Burger, C., Kuhlmann, A., Richard, P., Weinmann, J. Blockchain in the energy transition a survey among decision makers in the German energy industry. (2016)
4. J. Mattila, et al. Industrial blockchain platforms: An exercise in use case development in the energy industry. https://www.etla./julkaisut/industrial-blockchain-platforms-an-exercise-in-use-case-development-in-the-energy-industry/ (2016)

5. Stojkoska, BLR, Trivodaliev, KV. A review of Internet of Things for smart home:challenges and solutions. J Clean Prod, **140**, 1454-64.
6. C. Long, et. al: Peer-to-peer energy trading in a community microgrid. In: 2017 IEEE Power & Energy Society General Meeting, pp. 1–5. IEEE, Chicago (2017)
7. Christidis, K., Devetsikiotis, M.: Blockchains and Smart Contracts for the IoTs, IEEE Access, vol. 4, pp. 2292-2303 (2016)
8. Chitchyan, R., Murkin, J.: Review of Blockchain Technology and its Expectations: Case of the Energy Sector, arXiv preprint arXiv:1803.03567 (2018)
9. Andoni, M. et al.: Blockchain technology in the energy sector: A systematic review of challenges and opportunities. Renewable and Sustainable Energy Reviews (100), 143–174 (2019)
10. Mengelkamp, E., et al.: A blockchain-based smart grid: towards sustainable local energy markets. In: Computer Science-Research & Development, pp. 207-214, 2018
11. Murkin, J., Chitchyan, R., Byrne, A.: Enabling peer-to peer electricity trading. In: 4th International Conference on ICT for Sustainability. pp. 234-235 (2016)
12. Pop, C., et al.: Blockchain Based Decentralized Management of Demand Response Programs in Smart Energy Grids. Sensors. **18**(1), 162 (2018)
13. Oh, Se-Chang, et al.: Implementation of blockchain-based energy trading system. Asia Pacific Journal of Innovation and Entrepreneurship. **11**(3), 322–334 (2017)
14. Mihaylov, M., et al.: NRG-X-Change-A Novel Mechanism for Trading of Renewable Energy in Smart Grids. Smartgreens. 101-106 (2014)
15. Mannaro, K., Pinna, A., Marchesi, M.: Crypto-trading: Blockchain-oriented energy market. In: AEIT International Annual Conference, IEEE. pp. 1-5 (2017)
16. LeMahieu, C., RaiBlocks: A feeless distributed cryptocurrency network, (2017) https://raiblocks.net/ media/RaiBlocks Whitepaper English.pdf
17. Churyumov, A., Byteball: A decentralized system for storage and transfer of value (2015). https://byteball.org/ Byteball.pdf
18. IOTA: A cryptocurrency for the Internet of Things, www.iota.org
19. Buterin, V., Ethereum white paper. GitHub repository, 22–23 (2013).
20. Popov, S., The tangle, (2017). https://iota.org/ IOTA_Whitepaper.pdf
21. XJ Technologies Company Ltd, http://www.anylogic.com/
22. Macal, C. M., North, M. J.: Tutorial on agent-based modelling and simulation. Journal of Simulation **4**(3), 151-162 (2010)
23. Macal, C. M., North, M. J.:"Agent-based modeling and simulation: ABMS examples. In: Proc. of 40th Conf. on Winter Simulation. pp. 101–112 (2008)
24. Murkin, J., Chitchyan, R., Ferguson, D.: Goal-Based Automation of Peer-to-Peer Electricity Trading. In: From Science to Society, pp. 139-151 (2018)
25. Sheikh, Muhammad Ali.: Electricity trading among microgrids. Dissertation (2009)
26. Graphviz. Graphviz-Graph Visualization Software. https://www.graphviz.org.
27. UK Government Office for Science, Distributed ledger technology: beyond block chain. https://www.gov.uk/government/news/distributed-ledger-technology-beyond-block-chain,(2016)
28. Joi Ito, Our response to "A Cryptocurrency Without a Blockchain Has Been Built to Outperform Bitcoin", url=https://www.media.mit.edu/posts/iota-response/, accessed 15 March 2018
29. Janno Siim, DAG-Based Distributed Ledgers, 2018
30. Eric Wall, IOTA is centralized, url=https://medium.com/@ercwl/iota-is-centralized-6289246e7b4d, accessed 15 March 2018
31. IOTA Foundation, Coordinator, url = https://blog.iota.org/coordinator-part-1-the-path-to-coordicide-ee4148a8db08, accessed 15 March, 2018