



Implementing a blockchain-based local energy market: Insights on communication and scalability

Arne Meeuw^{a,*}, Sandro Schopfer^b, Anselma Wörner^b, Verena Tiefenbeck^{b,c}, Liliane Ableitner^b, Elgar Fleisch^{a,b}, Felix Wortmann^a

^a Universität St. Gallen, Switzerland

^b ETH Zürich, Switzerland

^c FAU Nürnberg, Germany

ARTICLE INFO

Keywords:

Local energy markets
Peer-to-peer trading
Blockchain managed microgrid
Smart grid communication
Network requirements

ABSTRACT

Peer-to-peer (P2P) energy markets are gaining interest in the energy sector as a means to increase the share of decentralised energy resources (DER), thus fostering a clean, resilient and decentralised supply of energy. Various reports have touted P2P energy markets as ideal use case for blockchain-technology, as it offers advantages such as fault-tolerant operation, trust delegation, immutability, transparency, resilience, and automation. However, relatively little is known about the influence of hardware and communication infrastructure limitations on blockchain systems in real-life applications. In this article, we demonstrate the implementation of a real-world blockchain managed microgrid in Walenstadt, Switzerland. The 37 participating households are equipped with 75 special smart-metres that include single board computers (SBC) that run their own, application-specific private blockchain. Using the field-test setup, we provide an empirical evaluation of the feasibility of a Byzantine fault tolerant blockchain system. Furthermore, we artificially throttle bandwidth between nodes to simulate how the bandwidth of communication infrastructure impacts its performance. We find that communication networks with a bandwidth smaller than 1000 kbit/s – which includes WPAN, LoRa, narrowband IoT, and narrowband PLC – lead to insufficient throughput of the operation of a blockchain-managed microgrid. While larger numbers of validators may provide higher decentralisation and fault-tolerant operation, they considerably reduce throughput. The results from the field-test in the Walenstadt microgrid show that the blockchain running on the smart-metre SBCs can provide a maximum throughput of 10 transactions per second. The blockchain throughput halts almost entirely if the system is run by more than 40 validators. Based on the field test, we provide simplified guidelines for utilities or grid operators interested in implementing local P2P markets based on BFT systems.

1. Introduction & motivation

For many decades, the power system in most countries followed a classic top-down architecture, in which large centralised power systems supplied a large number of consumers with electricity via a transmission and distribution grid. One promising path for the decarbonisation of the power system involves the deployment of renewable energy technologies, like photovoltaic solar power, which can be integrated as residential or commercial rooftop systems [1]. As more consumers transition into prosumers, the supply of electricity becomes increasingly decentralised, with many small prosumers who cover a part of their electricity demand with solar energy from their PV system. These prosumers still interact with the distribution grid on the one hand by injecting surplus solar energy (which they cannot consume themselves) into the distribution grid and on the other hand, by purchasing electricity from the grid when solar energy is not available (or in general,

when their demand exceeds their solar production). The growing share of Distributed Energy Resources (DERs) requires increased coordination efforts, due to intermittency and the bidirectional nature of the energy flows between many prosumers, consumers and the Distribution System Operator (DSO). The major coordination efforts include not only the balancing of demand with the intermittent supply of renewable energy, but also associating the energy flows with cash flows in order to properly bill consumers and reimburse producers.

Traditionally, grid control and coordination are conducted using a centralised management system, where a local or regional utility manages grid balancing, maintenance, customer billing, and applies regulations regarding remuneration of injected renewable energy. Consumers and prosumers are typically price-takers and have no say in how the price for injected energy is determined. In times of falling feed-in-tariffs [2], prosumers may face challenges to amortise investment

* Corresponding author.

E-mail address: arne.meeuw@unisg.ch (A. Meeuw).

costs [3] and the general investment dynamics of private households in roof-top solar energy may fall below the foreseen trajectory.

Local peer-to-peer (P2P) markets within a physical microgrid (all members of the local community are downstream a transformer station and share the same voltage level) may overcome those disadvantages by allowing consumers and prosumers to specify price limits and preferences, leading to potentially higher remuneration rates for prosumers and lower energy costs for consumers if they trade directly with each other [4]. A number of commercial projects have implemented this type of market between local producers and consumers. Notable examples are the British project *Piclo* [5] and the Dutch project *Vandebrom* [6] who both provide online portals to micro-source owners and consumers to match.

By using blockchain technology, P2P market places can be decentralised and enable consumers and prosumers to trade energy within a confined microgrid without relying on a central authority. The technology is expected to have disruptive potential and revolutionise local energy markets [7]. P2P markets are based on smart-metre data acquisition of consumption and production quantities. If the blockchain system is fully decentralised, the smart-metre must be extended to a computing device that can securely administer a blockchain account (private and public keys), issue transactions, or even act as a validator node. In addition to disintermediation, a blockchain-based microgrid system provides fault tolerant properties [8]. This means that a microgrid can still be operated if some validating nodes are malicious or offline.

However, the use case of blockchain-based Local Energy Markets (LEMs) is often discussed on a conceptual level [4] and there are still very few insights available from real-world field tests which utilise blockchain technology. In particular, it is unclear to how the requirements regarding data throughput are met in real-world conditions and to what extent the size or configuration of the P2P network play a role. In order to test the technical feasibility of blockchain managed microgrids, we have implemented a Local Energy Market with 37 participating households situated in the same physical microgrid based in Walenstadt, Switzerland. Prosumers and consumers are equipped with an extended smart-metre that hosts an application specific consortium-blockchain (a blockchain with a fixed set of smart-contracts needed for the application) and implements the LEM as a double-auction type order book.

In order to apply the concept of a P2P Local Energy Market to a real-world setting a number of design and engineering challenges must be solved. Although a high degree of decentralisation (i.e., high number of validators) can distribute the redundancy over a wider share of participants and correspondingly increase the resiliency of the system, a high share of validators can also increase the latency. The latency (time required for a transaction to be validated and included into a block) leads to lower transaction throughput due to the increased communication efforts between nodes. The maximum transaction throughput is of key interest for a LEM as it determines the minimum interval, under which demand and supply between the members of a microgrid can be matched. Transaction throughput and latency are major limitations of public and consortium blockchain networks. As a result, protocol developers often publish tests and present achieved transaction throughput in units of transactions per second (tps). As an example, public blockchains like Bitcoin or Ethereum achieve 7 and 15 transactions per second (tps), respectively, while consortium chains can scale easily beyond the order of 1000 tps [9]. However, these tests are usually conducted under idealised conditions to estimate an upper limit of the expected maximum transaction throughput. In addition, high performance servers may be used as validators for those tests with high-speed network connections between nodes. In local P2P networks, however, validator nodes may be highly constrained computing devices with limited computational resources and system memory. The data rate of the communication infrastructure used for smart-metring can be orders of magnitudes lower than those used in idealised tests.

There are currently no empirical benchmarks available to characterise the expected throughput and system latency for a given number

of validators and available communication infrastructure specifically for the application of a P2P LEM under constrained hardware environments. In this paper, we treat the data rate and number of validators as independent variables and observe the maximum transaction throughput and latency as dependent variables. The households used in the field-test have been equipped with 75 smart-metres, including SBCs. Therefore, the number of validators can be adjusted from 1 to 75.

This paper is structured as follows. In Section 2 we introduce academic literature and commercial projects related to the topics of this article (i.e., *Microgrids and local energy markets*, *Blockchain-based energy systems*, and *Smart grid applications and communication infrastructure*). Section 3 introduces the setup of the field-test in Walenstadt, its underlying infrastructure, and installed grid participant types. Furthermore, the section provides details on the local energy market application, processes, data structures, and a description of the underlying blockchain platform. In Section 4, we develop the applied test methodology based on tests conducted in related literature and present the results of these tests in Section 5. In section Section 6, we define a process of deriving the required communication infrastructure from the test results, benchmark the assumptions of the introduced field-test, and discuss the limitations of the pursued methodology. Finally, the findings of this study are summarised in Section 7.

2. Related work

2.1. Blockchain-based local energy markets (LEM)

Blockchain, in general, is an implementation of Distributed Ledger Technology (DLT) where a shared ledger (database) is kept among the participants of the network that aims to achieve constant agreement of all participants regarding its content [10]. A common property of all blockchains is that data is handled in blocks, which are chained after one another, and linked by the data hash of the previous block [11]. In Bitcoin, this data represents transactions of the underlying virtual currency, and balances of users are calculated from unspent transactions to their account [12]. In other implementations, such as Ethereum, block data represents a list of transactions between accounts holding embedded data and including mediating and governing programmes called smart contracts [13].

Since the first energy transaction over blockchain was performed by the Brooklyn Microgrid in Brooklyn, New York, in April 2016, the study of energy trading on the distribution level using a blockchain-based approach has taken off. The landscape of academic research on accounting schemes and technological assessment of blockchain-based LEM has been growing steadily since 2016. A significant number of businesses, such as Grid+, Power Ledger, and LO3 Energy have started offering new metring and billing solutions with blockchain-based technology and market schemes like P2P trading [14–16].

The mechanism to agreeing on the validity of an addition of a block to the chain is achieved by a system's consensus mechanism. The most prominent consensus mechanisms are Proof of Authority (PoA), Proof of Work (PoW), Proof of Stake (PoS), and Practical Byzantine Fault Tolerance (PBFT). Each of these protocols have advantages and disadvantages for their use in public and private blockchains and have to be selected for the respective use-case. Andoni et al. [17] provided a systematic overview of blockchain technology in the energy sector, and included detailed background information and details of various consensus mechanisms. The paper concludes that blockchain technology can clearly benefit energy system operations, markets and consumers, and offer novel solutions for empowering consumers and small renewable generators.

Mengelkamp et al. summarised the inner workings of the blockchain-based Brooklyn Microgrid and revealed a double-auction market with a single market clearing price [18]. The paper dissected the LEM into seven core components and elaborates on the consideration of each level within the Brooklyn Microgrid. Offering a deeper insight into

the actual transactional and functional structure of their implemented system, Sikorski et al. [19] employed a P2P market between two energy producing and one consuming machines in the chemical industry using a round-robin scheme for validator selection. Three virtual machines run the blockchain, and consumption and production data is simulated by the Aspen Plus modelling software [20]. The article did not present results from the energy exchange or any further analysis of the built system, but provided a good introduction to blockchain technology. A list of projects for microgrids utilising blockchain technology, both research and commercial, is presented by Goranovic et al.. The authors noted that decentralised systems require increased communication speeds [21].

Several publications introduce blockchain-based transaction schemes and LEM models, to overcome centralised markets with a traditional pricing model [22] or a continuous double auction [23], to manage demand response and activate financial settlement for flexibility providers [24], and to mitigate participant errors by employing outage detection [25]. However, none of these consider the underlying infrastructure nor the communication requirements.

To guarantee fair rules for all participants, such as producers, consumers, and storage owners, the market application can be run either by a trusted central party, by a subset of trusted participants, or all participants. Since the dependency on a single central authority is neither a decentralised system nor a fault-tolerant system, the latter two are the preferable options within this context.

2.2. Communication infrastructure for smart grid applications

The vision of the smart grid includes the holistic integration of information from the power system infrastructure. This enables renewable energy systems, consumers, and power plants to seize the full potential of the connected energy resources and maximise the efficient usage of the grid infrastructure. The integration of applications spans from the energy suppliers, over the transmission grid, to the distribution operators, and finally to the retail customer [26]. These services and applications include Advanced Metering Infrastructure (AMI), real-time pricing (RTP), Demand Side Management (DSM), Demand Response (DR), distributed generation, storage, and others. The requirements for these applications are numerous and include the communication between endpoints, measurement equipment, and control infrastructure.

A local energy market combines smart grid applications, such as advanced metering, real-time pricing, and (optionally) demand-side management into a single application. Smart metering involves communication channels from the metering data management system to the distributed measurement points in consumer and prosumer households.

Liang et al. [27] introduced a demand-side energy management scheme in residential smart grids and note that most existing studies assume perfect two-way communication, which is unrealistic for practical applications.

Ancillotti et al. [28] defined a number of quantitative and qualitative communication requirements for smart grid applications. The most important quantitative requirements are the *data rate*, *latency*, and *reliability* of a system. The paper mentions that the microgrid concept is not new and industrial microgrids are a good example realised in practice and in research (see [19]). Kuzlu et al. gave a comprehensive overview of communication requirements for major smart grid applications [29]. In regard to the applications within the Neighbourhood Area Network (NAN), scheduled metre readings by AMI, time-of-use (TOU) and RTP, DR, and Distribution Automation (DA) are listed with typical message data sizes and maximum tolerable system latency. Typical AMI payload sizes are quoted between 100–200 bytes excluding transport protocol overhead by [30]. A confirmation of these values for smart grid applications, such as on-demand and multi-interval metre reading is given by [31], which specify typical metre reading messages being in the range of 100–200 bytes and required data rates of about 100 kbit/s per device.

Zia et al. [32] emphasised the importance of communication infrastructure to share information for optimised operation when dealing with dispersed generation in microgrids. The authors noted the necessity for reduction of installation costs by selecting suitable data communication technology for short and long distance applications, and suggest that wireless technologies are more suitable due to their lower deployment cost.

Another work, focusing on a specific communication technology and its suitability for smart grid services, is presented by Li et al. [33]. The paper lists a variety of communication technologies and compares the capabilities of the Narrowband Internet of Things (NB-IoT) standard with quantitative requirements of smart grid services similar to [28]. The authors mention data rate requirements for Demand Response Management (DRM) related services around 14–100 kbit/s, while RTP and TOU pricing schemes may require additional 100 kbit/s.

While in some regions, coaxial cable infrastructure is widely available due to television deployment, other regions have already invested in fibre optic cable infrastructure. Technologies, such as Digital Subscriber Line (DSL), offer broadband data rates over existing telephone lines, but are usually used for consumer connection purposes. Power Line Communication (PLC) can be used on existing power lines and provide an independent communication channel to transmit data. While bridging the gap between communication endpoints over the air, wireless technologies usually have the advantage of lower installation costs. When depending on deployed infrastructure, such as cellular, service charges may apply according to the required bandwidth and volume. Open and freely operable technologies, like Wireless Personal Area Network (WPAN) and Low Power Wireless Personal Area Network (LP-WPAN), offer good ranges but are very limited in their maximum bandwidth capabilities. Even though open and crowd-sourced network initiatives, such as The Things Network (TTN) [34], exist for WPAN technologies like Long Range Wide Area Network (LoRaWAN), they are still dependent on relay points, which are also commercially distributed by communication providers. In addition to data rate and latency limitations, the reach of wireless technologies into buildings and lower levels, where metering infrastructure is often deployed, is limited.

2.3. Review of throughput tests of blockchain platforms

Buchman [35], presented a variety of tests, in which the transaction throughput limits of the Tendermint consensus mechanism are quantified. The tests were run with a *nil* application in order to disable the Tendermint Socket Protocol, as well as the mempool. Transactions had a size of 250 byte/transaction and were preloaded on the validators. The used machines were high performance Amazon EC2 instances with 4 & 32 GB of Random Access Memory (RAM) spanned around the globe and each validator was directly connected with one another. The tests offer a good methodology, but are limited in their validity concerning real world applications, since they show the absolute limits of the Tendermint consensus algorithm. Under those idealised conditions, Buchman [35] showed that transaction throughput rates well beyond 1000 tps can be reached.

Han et al. [36], provided a series of experiments which evaluated the latency and throughput of the Ripple blockchain network, as well as the Practical Byzantine Fault Tolerance (PBFT) based Hyperledger Fabric (HLF) consensus algorithm. In both systems, transactions were sent as a Javascript Object Notation (JSON) object via a Remote Procedure Call (RPC) interface over Hypertext Transfer Protocol (HTTP) and represent a value transfer of the amount 1 monetary unit between two accounts. In order to realise value transfer in HLF, a chaincode implementing a simple money transfer application was deployed on the cluster.

The authors claim to evaluate blockchains for the Internet of Things (IoT), but employed three high performance computer instances with multicore processors and 12 GB RAM, without providing any means of limiting the computing performance. The validators are realised as

virtualised instances running in docker containers and details regarding network interconnection or limitations are not provided.

Blom and Fahramand [37] provided a study on the scalability of a blockchain-based LEM. They used a private deployment of the Ethereum blockchain platform and measured the amount of information the platform needed to process a day-ahead market and a real-time energy market. Their evaluation provided requirements for minimum transaction throughput for a base case scenario of a 5 min real-time energy market with 600 participants, and included a sensitivity analysis of the system for different trading frequencies, as well as participant numbers.

BLOCKBENCH, introduced by Dinh et al. is an evaluation framework for analysing private blockchains [9]. It is designed to integrate and benchmark the latency, throughput, scalability, and fault-tolerance of any private blockchain. The article provides a variety of tests of the Ethereum, Parity, and Hyperledger blockchains and measures maximum performances, using a 48-node commodity cluster. Every node of the cluster has a E5-1650 3.5 GHz CPU, and 32 GB RAM.

2.4. Summary of related work and research focus

To sum up, the current academic literature provides bandwidth requirements for AMI applications in smart-grids. However, those benchmarks are in all cases reported for centralised systems that pull information from deployed AMI devices. However, blockchain based systems operate in a decentralised fashion and require therefore higher bandwidth to operate due to synchronisation processes between the distributed nodes. The academic literature provides few benchmarks bandwidth requirements of decentralised architectures for AMI applications in the energy sector.

The transaction throughput rate (measured in tps) is a core characteristic of various blockchain platforms. The throughput rate depends on a large parameter space (like consensus mechanism employed, number of validators, used hardware, complexity of the application that runs on top of the blockchain etc.). Due to the recent advent of blockchains, little is known about throughput rates under constrained environments faced in practical applications like LEMs. The main constraints in LEMs are low computing power of smart-metre devices and potentially limited bandwidth of the underlying communication networks. The present article studies the effect of reduced computing-power, degree of decentralisation and bandwidth on throughput for LEMs. We study the throughput rates in a real-world environment, which provides realistic and field-tested design guidelines that support utilities and service providers when designing blockchain-managed LEMs.

3. Application and system design

In this section, we introduce the deployed field test setup, devices and configuration of the Walenstadt microgrid, followed by a detailed description of the implemented market application, as well as an outline of the blockchain platform design.

3.1. Walenstadt community microgrid

The Walenstadt community microgrid consists of 37 households in Walenstadt, Switzerland. The community includes 25 single family prosumer households, eight of which have local storage installed, as well as two apartment buildings with installed PV, of which one has a battery. In addition to the prosumer households, the microgrid includes two single family consumer households and one elderly home. The total installed PV power is 287 kW, while the yearly consumption of the participants is about 460,000 kWh, of which the elderly home takes a large share of about 200,000 kWh. The total installed storage capacity is around 80 kWh.

Within the Walenstadt microgrid, every device which participates in the LEM, is represented by its own computing device. Production,

Table 1
Building and metre types of the field test setup.

Building	Consumption	Collective consumption	Production	Storage
Apartment Building 1	7	1	1	1
Apartment Building 2	2	1	1	0
8 × Prosumers with storage	1	0	1	1
17 × Prosumers	1	0	1	0
3 × Consumers	1	0	0	0
Sum	37	2	27	9

consumption, and storage values are measured separately from each other, which yields a total install base of 75 Raspberry Pi SBCs. These devices are running the blockchain platform and market application as a decentralised system. Fig. 1 shows building types and installed devices of the Walenstadt microgrid setup. Table 1 lists the number of devices and types of the measured consumption or production. Each of the Raspberry Pis runs on a quad-core ARMv8 Central Processing Unit (CPU) clocked at 1.2 GHz and 1 GB of RAM. Each device is equipped with a SmartPi Hardware Attached on Top (HAT), which enables it to measure voltage and current of three phases [38]. Fig. 2 shows a picture of one of the participating apartment buildings, as well as two SmartPi devices installed in a prosumer household circuit breaker box.

The devices are interconnected in a Virtual Private Network (VPN) with internet access via a cable modem available on the premises. Depending on a building's configuration, multiple devices are to be connected in a Local Area Network (LAN), via a common cable modem by use of the internal switch of the modem, or a 100 Mbit Ethernet switch. In two houses, an inconvenient location of the circuit breaker made it necessary to deploy local PLC. A measurement of the available data rate using the *iperf* utility yielded 5550 collected data points, which includes the maximum data rate over ten seconds between each of the 75 utilised devices. Each data point represents a connection between two devices, as all the devices act as both a server and a client. Fig. 3 shows the results and provides an overview of the interconnectivity of the deployed devices in the field test. The results of the interconnectivity test show an average data rate of 1.7–2.6 Mbit/s when utilising the coax cable connection, and a data rate of 93.8–94.5 Mbit/s, when devices are connected to a LAN.

3.2. Market application

The Walenstadt microgrid enables participants to trade local energy with each other when available, and sell or buy residual energy to and from the grid. Participants are provided with a web-based dashboard application, which they can use to define their buy and sell price preferences, and oversee their consumption, production, as well as the realised market prices over time. To enable this functionality, we are running a LEM application, which is entirely executed as a consensus afflicted application on a blockchain platform. In this section, we introduce the functions and data structures of the market application and go into further detail of the blockchain platform in Section 3.3.

Fig. 4 shows an overview of the phases of the market application and its collected and calculated data structures. Market data collection starts by receiving transactions from the agent, which contain a specific payload describing an order. Every *clearing interval* – a common parameter of a real time energy market – the collected orders are matched according to the underlying market mechanism, which yields a set of trades between the participants of the system. The market mechanism is based on a double-auction mechanism [18,23]. Every *settlement interval* – another predefined parameter – accumulated trades are summed up and yield a list of settlements, which represent the final financial flows between participants.



Fig. 1. Setup of the Walenstadt microgrid. The district consists of 37 participants with varying configurations. Two apartment buildings, eight prosumers with local storage, seventeen prosumers, two consumer households, as well as an elderly home as a large consumer.



Fig. 2. Photos from the Walenstadt microgrid. The left picture shows the large prosumer apartment building with seven apartment units, a 35 kW PV system, and a 23 kWh battery. The right picture shows two SmartPi devices and attached current transformers (in red) installed in a circuit breaker box in a single family prosumer building. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Distribution of the data rate benchmark results between participants. Each data point represents a unidirectional data rate measurement between two devices, resulting in $75 \cdot 74 = 5550$ data points. The majority of connections are via coax cable and reach data rates between 1.7–2.6 Mbit/s, while locally interconnected devices reach between 93.8–94.5 Mbit/s.




Fig. 4. Flow of information through the market application. A transaction containing the order induces the bidding process, which posts the corresponding order to the list of orders. Every 15 min, the clearing process matches the existing orders into trades between participants. Every 24 h, the settling process aggregates the accumulated trades into settlements for payment distribution.

Bidding. During this phase, the market is open to collect orders from the connected agents and fill the order book. This phase, the *bidding period*, is defined by the clearing interval of the LEM, and is open between the time of the last clearing until the next time interval, minus a security margin, called the closing time. The structure of a transaction is shown in Listing 1. The *payload* is used to add application-specific values to a transaction, such as posting an order to the market. Listing 2 shows the structure of a transaction payload for issuing an order on the market. Necessary values for the interactive market include the buy and sell price, as well as the amount of energy to be traded. The field `gridUnits` quantifies energy to be sold or bought exclusively to and

from the grid. The field is used in case an agent was offline during a previous bidding period and has missed the chance to trade its energy on the real-time LEM. When receiving a transaction, the blockchain platform application logic checks the signature of the transaction and forwards it to the market application. The market application validates the transaction’s payload, checks the participant database for affiliated devices, and (if successful) adds the corresponding order to its stash of orders. Orders are derived from the data available from both the transaction (`senderAddress`), as well as the payload fields (`buyPrice`, `sellPrice`, `units`), and can be seen in Listing 3.

```

1 TxStruct {
2   senderPubKey: string;
3   receiverAddress: string;
4   senderAddress : string;
5   type: string;
6   payload : string;
7   signature : string;
8   value : number;
9   nonce : number;
10  hash : string;
11 }

```

Listing 1: Structure of a blockchain transaction. The field values include the addresses of the sender and receiver ('senderAddress', 'receiverAddress'), the sender's public key ('senderPubKey'), the transaction value, a field for a payload as well as the transaction specific 'nonce' and 'type'.

```

1 OrderPayload {
2   buyPrice: number;
3   sellPrice: number;
4   units: number;
5   gridUnits: number;
6 }
7 %

```

Listing 2: Structure of a transaction payload to issue an order. The field values include the buy price and sell price, and the amount of units. The field `gridUnits` denominates energy amounts not traded on the local market.

```

1 OrderStruct {
2   bidder: string;
3   buyPrice: number;
4   sellPrice: number;
5   units: number;
6   nonce: number;
7   timestamp: number;
8   hash: string;
9 }

```

Listing 3: Structure of an order. The fields include the address of the bidder, the proposed buy- and sell price, the amount of units, as well as a timestamp, the underlying transaction's nonce and hash.

```

1 TradeStruct {
2   buyer: string;
3   seller: string;
4   unitPrice: number;
5   units: number;
6 }

```

Listing 4: Structure of a trade between two participants. The fields include the addresses of the buyer and the seller, the matched price and amount of traded energy.

Clearing. The clearing process is triggered every clearing interval by the application logic. By comparing the current timestamp (e.g. the timestamp of the last finalised block) with the timestamp of the last clearing event, the application executes the clearing process. The clearing is defined by the market logic and is subject to the individual bidding language and market mechanics of the platform. Clearing trades essentially transforms the orders collected during the *bidding period* into a list of trades among the participants. Fig. 5 shows the types and number of trades, as well as the trading volume over the course of a sunny day in February, 2019. The more local energy is available, the more trades are matched between local participants, as a single consumer can match its energy demand from one or more supplying prosumers. Residual supply or demand is always covered by the grid. Listing 4 shows the structure of a trade, which consists of the buyer's and seller's address, as well as the agreed price and matched units. In our case, this step occurs every 15 min as shown in Fig. 4.

Settling. The last step of the market application logic is settling matched trades into settlements. Settling is called every 24 h and aggregates the accumulated trades over the past settlement period into

a list of settlements, which are used for automated payments or further billing purposes.

3.3. Blockchain platform

As mentioned in Section 2.1, blockchains store their data entries in blocks, interlinked by a checksum of the contained data (hash). There are many consensus protocols being used to reach agreement between the participants over the current state of a blockchain system. The most common ones are Proof of Work (PoW), Proof of Stake (PoS), and Byzantine Fault Tolerance (BFT).

Proof of Work (PoW) is the most prominent consensus protocol, as it has been used to secure the global Bitcoin network for more than 10 years. In proof of work, so called miners compete for a valid solution to a cryptographic puzzle, which can only be solved by brute force. The difficulty of this puzzle is constantly adjusted (the number of most significant bits in the calculated hash being zero). By means of this adjustment the computing power of the whole network is able to find one solution to the puzzle every ten minutes. Once a miner has found a valid hash, it propagates the new block containing the required data to the other miners in the network. Each miner verifies the validity of the proposed block and restarts its puzzle solving based on this new block. PoW based blockchains offer stable block intervals, as the network periodically adapts to the difficulty of the cryptographic puzzle to match the network's total computing power.

Proof of Stake (PoS) requires the validating nodes to put a portion of their economic stake on the platform as collateral, instead of an investment into computational power. The possibility of losing stake in the system because of a dishonest proposal provides an incentive for the validators to play by the rules. There are two major schools of thought in the design of such algorithms:

1. The *chain-based* PoS mimics the mechanics of PoW and seeks to replicate the pseudorandom selection of validating nodes.
2. The Byzantine Fault Tolerance (BFT) based proof of stake reaches back to over 30 year old research on the Byzantine Generals Problem and protocols to solve it [39]. A group of validators try to reach consensus on the execution order and results (state) of transactions while a subset may behave arbitrarily (Byzantine) faulty. This mechanism employs voting and timing schemes to identify these faulty nodes and manages to tolerate up to f faulty nodes, as long as $3f + 1$ nodes behave correctly. This consensus process requires many exchanged messages between validators to finalise. The block creation intervals are based on defined response timeouts of the validators, as well as any delays within the P2P messaging.

The Tendermint consensus protocol provides a framework to keep replicated state machines in synchronisation between an arbitrary set of computing nodes, and is based on BFT. We use the Tendermint protocol for the blockchain platform, which enables the market application to be executed as a smart contract [40]. We chose the Tendermint consensus protocol, as it provides a basis for building performant private blockchain systems between a number of known participants. Next to Hyperledger Fabric, it is one of the most mature, stable, and useable consensus protocols. This stability and robustness made it the candidate of choice to build the decentralised system at hand. The flexible Tendermint eco-system offers a range of frameworks to write an application with, such as the Cosmos SDK for applications written in Go, and the Lotion framework, for applications written in Javascript/Typescript. The interoperability of the Cosmos Hub to side-chains, as well as the PoS consensus mechanism securing the main chain, allows for future interconnection of many energy markets. Using a consensus mechanism such as PoW for a local energy market goes against the principle of reducing the overall energy consumption. We have built the market application compatible with the Tendermint specification and can scale the degree of decentralisation from one node to any number of nodes,

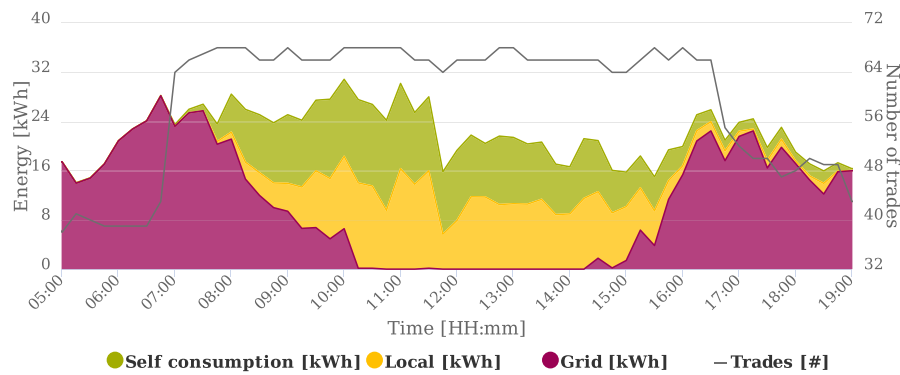


Fig. 5. Trading volumes between participants and the grid in Walenstadt over the course of February 14th, 2019. Red shows the volumes traded with the grid, green shows the self consumption of participants, and yellow shows the volumes traded between participants. The grey line shows the number of trades in each 15 min time slot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

as long as the available computing power and communication between validators is sufficient. To achieve consensus over the creation of a new block, data needs to be propagated to the validators of the system. Transactions are checked by a validator according to the application logic when they are received and added to an ordered list, called the *mempool*. This list is then broadcasted to the other validators [35]. If the result of the transaction-induced state transitions match for at least 2/3 of the system's validators, consensus is reached and a new block is created. Blocks include transaction data as well as the state of the resulting application. Through this process, we ensure that the application is run in its intended form, as long as less than 1/3 of the validators are not faulty or malicious (i.e., Byzantine) [39].

The market application, described in the previous section, is a transaction based finite state machine, which maintains its state as a JavaScript object. The data held in the state includes participant account information, collected buy and sell orders, matched trades, as well as settlements. The application processes information (orders, trades, settlements), which are financially binding, and thus, the execution of the bidding, matching, and settlement processes must be agreed upon by the participants of the community. The system consists of three different participant types: The core of the platform is built by its validator nodes, which are represented by the prosumers and the utility company (energy producing entities). Consumer nodes are clients of the prosumers and do not propose new blocks on their own, but rather read and send transactions by relaying their communication to the blockchain over the validators.

The Raspberry Pi SBCs of the Walenstadt microgrid run software depending on their role in the system. Devices representing PV systems run the blockchain platform and LEM application as validators, in addition to the agent software. The *agent* software runs on every device to obtain measurements and translate user preferences into orders.

The bidding process of the LEM is invoked by incoming transactions to the blockchain platform, while clearing and settling are time-triggered processes executed by every validator. Clearing and settling are not based on transactions, but still alter the state of the system (e.g. take the collected orders, match into trades and remove them from the stash) for regular time intervals, which are executed less often than the processing of an incoming transaction (15 min and 24 h). The bidding step processes a single transaction for each of the participants within each bidding period. This makes it the most communication intensive process of the application. Hence, we focus on the speed limits of the bidding process and measure its maximum transaction throughput.

Reference application: cointest. In addition to the introduced LEM application, we run a similar example to the value transfer example used by [36]. In order to evaluate the maximum transaction throughput of the Raspberry Pi SBC devices, an application, referred to as *cointest*, is

run which includes a signature validity check of the incoming transaction and a manipulation of a key-value store to determine an exchange of funds between accounts induced by a transaction.

The complexity and therefore the computational expenditure of the market application's bidding phase is far higher than the value transfer logic of the cointest application. While both applications check a transaction's signature validity, the cointest application finishes with the subtraction of the transaction value from sender's account and the addition of the transaction value to the receiver's account. The market application, however, further checks the validity of the transaction payload, checks for affiliated participants to allocate shares, which involves multiple iterations over participant lists, and finishes by adding a bid to the order list.

4. Test methodology

Overview of conducted tests. The focus of this paper is to quantify how bandwidth and the degree of decentralisation (number of validators) affects throughput rate and transaction latency using constrained hardware (smart-metre with integrated SBC) as deployed in the Walenstadt field test. To fully characterise the blockchain performance deployed in the field, we run two distinct tests, which are illustrated in Fig. 6

Test 1 aims at quantifying the impact of the market-application and the required computing power on the throughput rates and transaction latency. Comparing the throughput rate of the P2P market against the cointest application (an application that only confirms transactions, i.e. a simple ledger that allows the devices to send transactions without the complexity of a P2P market application) allows us to infer the maximum throughput limit for both the cointest and P2P energy market application with a single validator node. This is an important benchmark, because existing studies [9,37] test maximum throughput limit using high performance cloud computing nodes rather than restricted hardware as deployed in field test described in the present article.

On the other hand, test 2 involves the entire market application with various numbers of validators at different bandwidth rates. For a given bandwidth, the throughput rates will therefore always be smaller or equal than the throughput rates of test 1 due to the higher number of validators.

Test parameters and variables. The relevant data points have been derived from the quantitative requirements of the communication infrastructure of smart grid applications. These Requirements include data rate in kbit/s and latency in s. As blockchain systems are based on transactions which alter the state of the system, we measure the throughput in successful *transactions per second*. Our testing environment parameters include values specific to the different parts of the system; namely the *loadtest*, *requests*, *data rates*, *validators*, and *iterations* (see Table 2).

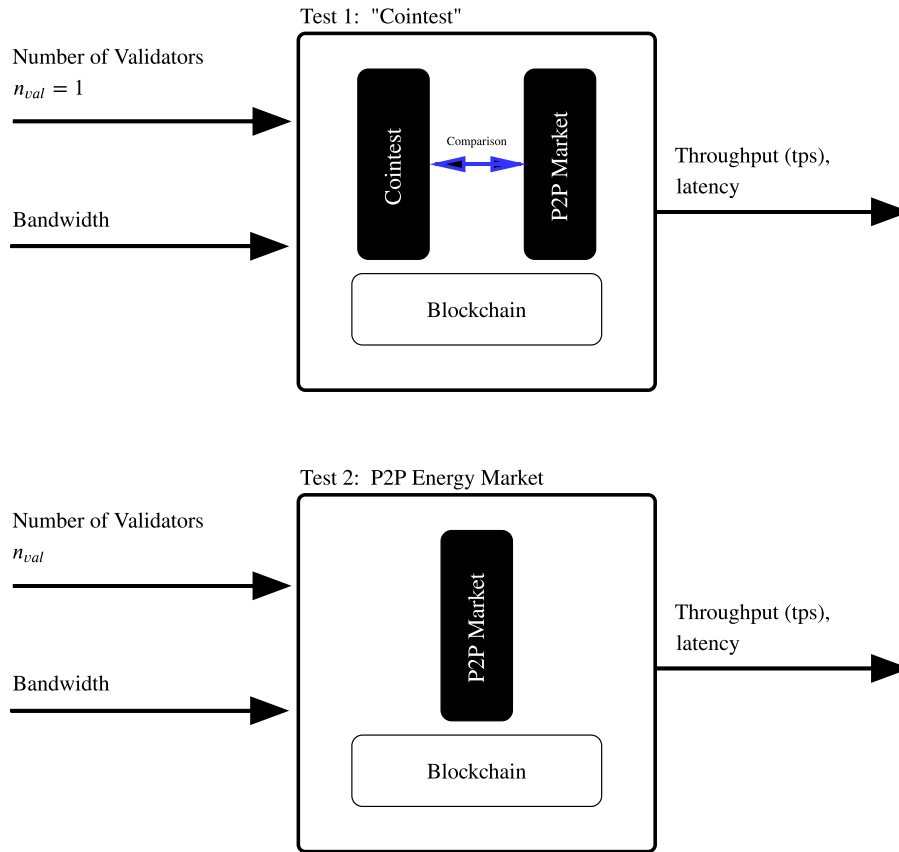


Fig. 6. Overview of conducted tests: Test 1 aims at providing the upper throughput limit for the chosen hardware and at different bandwidth. Test 2, on the other hand, provides throughput rates for different numbers of validators and bandwidths for the P2P energy market application.

Table 2
Parameters and experimental settings.

Parameter	Description	Setting
Loadtest		
URL	Host Name or IP Address	"privateIP:46659"
concurrency	Number of simultaneous requests	1
maxSeconds	Duration of test	180
maxRequests	Number of total successful requests	22–450
Requests		
startRps	Start value of requests per sec.	0.25
maxRps	End value of requests per sec.	5
stepSize	Increment step of requests per sec.	0.25
Data rates		
startIndex	Start of data rate set index	0
maxIndex	Maximum data rate set index	16
set	List of data rates in kbit/s	50–16 000
Validators		
startIndex	Start of validator set index	0
maxIndex	Maximum validator set index	11
set	List of validators	1–64
Iterations		
startIndex	Start of iterations	0
maxIndex	End of iterations	4

The *loadtest* module provides a framework to conduct HTTP throughput tests. The module sends a defined amount of requests per second and returns statistics of successful requests, latency, and errors. There are several different options which the range of requests per second can be tested with, such as the maximum number of requests,

the maximum duration of a test, and the target host. Data rates are limited to a certain set of values, defined by the *datarates* options, which are selected to represent a variety of available communication protocols. We artificially limit the maximum data rate on each device using the *Wonder Shaper* script. *Wonder Shaper* has been utilised before by [41], in order to limit bandwidth, measure the minimum required bandwidths, and to emulate different network conditions on fog devices, as was done in [42]. The *validators* options define the range of validators the system is tested with. The developed test framework connects to the desired number of validator devices, and launches the blockchain-platform with the desired parameters. These include validator key-pairs, genesis files. The *iterations* options define repetitions of the conducted tests. The tests are run multiple times (five times) and the average values are then used to benchmark the performance.

We run the tests against the `/broadcast_tx_commit` endpoint of the Tendermint consensus engine in order to test the throughput and latency of the *final* transactions, which have been processed and accepted by the market application. This endpoint takes transaction data encoded in hexadecimal format and returns the request as soon as the transaction is committed into a block. For details about the Tendermint RPC, please refer to the Tendermint documentation available at [43].

5. Results

In this section, we provide the results performed following the introduced methodology. First, we present a maximum transaction throughput test to show up the computational limits of the utilised hardware. Second, we show the results of the tests performed with validator systems of low (i.e. 1–8 validators), medium (12–40 validators), and high (48–64 validators) degrees of decentralisation. We ran the

tests with artificially limited data rates from 50–16 000 kbit/s, as well as no limit. The results include the maximum achievable transaction throughput and the latency of final transactions.

The *field test* setup, includes all devices of the Walestadt microgrid field test, as well as a server connected to the VPN functioning as the testing client.

We ran tests of 1, 2, 4, 8, and 12 validators at 17 data rate limits, and 16, 24, 32, 40, 48, 56, and 64 validators at 6 data rate limits. The combinations amount to 127 configurations of validators and data rate limits. The results are averaged over five repetitions of each test-configuration.

5.1. Test 1: Computational limits of the Raspberry Pi SBC

The tested hardware platform induces a set of limitations, due to its computational capabilities. In order to separate computational limits from the device and limits induced from the communication infrastructure, we test the cointest application, as well as the market application on a single Raspberry Pi.

In order to estimate the maximum transaction throughputs at a certain data rate, we averaged 1000 generated transactions and found an average encoded transaction size of 352 bytes for a cointest transaction, and 454 bytes for an order transaction (for the market application). The overhead caused by the underlying transport protocols, namely Transport Control Protocol (TCP) and HTTP, yields request sizes of 859 bytes, and 1061 bytes respectively.

Fig. 7 shows the maximum average transaction throughput of the cointest, as well as the market application processed on one validator at data rate limitations ranging from 50 kbit/s to no limit (i.e., 100 Mbit LAN). The dashed lines in the graph show the theoretical maximum throughput values for the two applications derived from their request size. The theoretical limits do not account for response traffic and are an upper bound of maximum transaction throughput at each data rate limit. The cointest application is clearly limited by the maximum data rate of the network and follows the corresponding theoretical maximum values with an offset factor of 0.65x. The throughput peaks at approximately 62 tps and reaches a plateau at 750 kbit/s. At 750 kbit/s, the theoretical maximum for the cointest application is >109 tps. The LEM application falls in line with about 0.7x of the theoretical maximum and reaches a plateau of approx. 10.5 tps at 200 kbit/s. The theoretical maximum throughput for market application transactions at 200 kbit/s is >23 tps.

The discrepancy between the limitations of the data rate and the measured values is caused due to the increased complexity of the processing for the market application versus the processing for the cointest example (as detailed in Section 3.3). In essence, this test shows an expected maximum transaction throughput of around 10.5 transactions per second (tps) for the market application, due to the limited single-core performance of the utilised Raspberry Pi SBC. The described tests provide insights into the performance of single validator performance measures, however provide no insights into the performance of multi-validator systems. In addition, BFT systems often require as many as $\mathcal{O}(n_{val}^2)$ messages per block, where n_{val} is the number of validators. The following tests provide further insights of the data rate requirements of a multi-validator system, with increasing degrees of decentralisation.

5.2. Test 2: Transaction throughput limits of P2P energy market application

Low to medium degrees of decentralisation (1–12 validators) at low data rates. Fig. 8 shows a heatmap of the throughput results from 1–12 validators at data rates from 50 kbit/s to 750 kbit/s. The tests from the upper data rate limits of 1000 kbit/s–no-limit are provided in the analysis from low to high degrees of decentralisation for better comparison.

The results show that a single validator (i.e. centralised) systems offers the highest transaction throughput, peaking at the found limit

of approx. 10.5 tps. This peak is reached at a data rate as low as 150 kbit/s. With increasing numbers of validators, the transaction throughput decreases due to increased coordination between validators, causing higher demand of network communication, as well as computational overhead. While lower degrees of decentralisation (1–4 validators) achieve medium to high transaction throughputs of 5–10 tps at 350 kbit/s, medium degrees of decentralisation (8 and 12 validators) do not establish a stable network operation below data rates of 250 kbit/s. A minimum requirement of 250 kbit/s suggests that constrained communication infrastructures, such as WPAN, narrowband PLC, or NB-IoT, can only support feasible operation for the lowest degrees of decentralisation (1–4 validators). In accordance with the communication bound limits of BFT systems, the data rate limit has an increasing impact with higher degrees of decentralisation. This is shown by the increasing throughput numbers at higher data rate limits.

Based on these results, we limit the amount of requests per second to 0.5–5 tps and advance in 0.5 tps steps when testing the field test setup with 16–32 validators. Furthermore, we reduce the requests per second to 0.25–1.25 tps and advance in 0.25 tps steps for decentralisation degrees of 40–64 validators. The results show that data rate limits under 1000 kbit/s are barely enough to enable stable operation for twelve validators. Therefore, we start the data rate limits for the following tests at this minimum.

Low to high degrees of decentralisation (1–64 validators) at high data rates.

Fig. 9 shows a heatmap of the throughput results from 1–64 validators at data rates from 1000 kbit/s to no limit (approx. 2.2 Mbit/s as measured in Section 3.1). The axes of Fig. 9 are transposed for readability. The results confirm that the maximum transaction throughput of a single validator system is approximately 10.5 transactions per second, independent of the available data rate (>1000 kbit/s). All of the tested configurations offer stable network operation at a minimum data rate of 1000 kbit/s, however, the configurations show a varying impact of the available data rate on the maximum transaction throughput.

While lower degrees of decentralisation (1–8 validators) show a high impact of the data rate at lower values (Fig. 8), the impact becomes less at higher data rates, suggesting that the computational capabilities of the utilised hardware are saturated.

For medium degrees of decentralisation (12–40 validators), the results show a stronger impact of data rate limitations on the maximum transaction throughput, increasing with the available data rates. While a 32 validator system requires at least 2000 kbit/s to process over 1 tps, a system with 12 validators already processes over 4 tps at this data rate.

For high degrees of decentralisation (48–64 validators), the data rate limitations do not seem to impact the transaction throughput, which stays at a constant 0.4 tps. This is due to the coordination between validators being computationally so intensive for the Raspberry Pi SBC, that higher transaction throughputs cannot be reached.

These tests show that medium to high degrees of decentralisation (12–64 validators) require a relatively high data rate to provide stable operation. The required data rates of over 1000 kbit/s, cannot be provided by a variety of communication infrastructures. WPAN technologies, such as ZigBee, LoRaWAN, and SigFox only provide data rates up to 250 kbit/s. Neither narrowband PLC communication, nor cellular technologies such as GSM or NB-IoT can provide data rates over 1000 kbit/s, and are thus not suitable for medium to high degrees of decentralisation.

5.3. Test 2: Transaction latency at maximum transaction throughput of P2P energy market application

Latency represents an important quantitative requirement for smart grid applications and infrastructure. The latency of a system defines how fast the connected devices are able to react to changes of the system. In case of a LEM, these limits are within the range of 5–60 s (confer with Section 2.2). However, other smart grid applications,

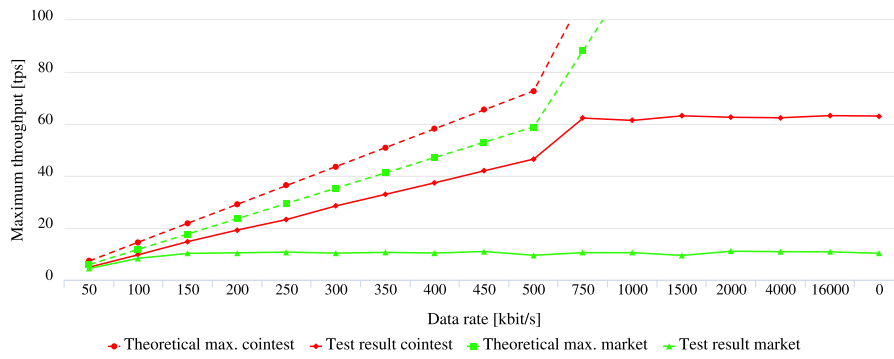


Fig. 7. Maximum transactions per second at different data rates. The test is done with the coinbase example to provide an estimation of the maximum achievable throughput using one validator.




Fig. 8. Maximum average transactions per second (tps) at data rates from 50–750 kbit/s and decentralisation degrees from 1 to 12 validators. The tests show the average values from five repetitions.

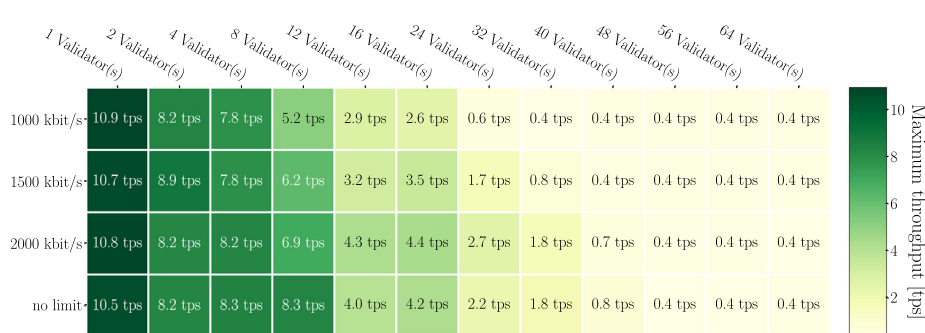


Fig. 9. Maximum average transactions per second at data rates from 1000 kbit/s to no limit, and decentralisation degrees from 1–64 validators. The results show the average values from five repetitions.

such as control of Distributed Energy Resource (DER) and Distribution Automation (DA), require lower latencies of 300 ms to 2 s [33].

In order to provide a guideline of the extensibility of the introduced blockchain platform for other smart grid applications, we provide the resulting latencies of confirmed transactions, in addition to the maximum transaction throughput. The results in Figs. 10 and 11 show the mean latencies of the transaction requests at the maximum throughput values presented in Section 5.2. As explained in Section 4, the conducted tests measure the transaction throughput and latency of final transactions. The latency of the system defines the amount of time necessary to finalise a new state and make it available to its clients. In blockchain-based systems, this time is defined by the block interval, which – in BFT based protocols such as Tendermint – depends on the consensus process between the blockchain platform’s validators.

Fig. 10 shows the resulting heatmap of the mean latencies of final transactions, for configurations of 1–12 validators and data rates from 50 kbit/s to 750 kbit/s. The latencies range from 9.8 s to 1.2 s and

generally decrease with increasing data rate limits. The results show that the latencies increase significantly with increasing numbers of validators, which falls in line with the communication bound limits of BFT systems.

Analogous to the results presented before, Fig. 11 shows a heatmap of the latencies for degrees of decentralisation of 1–64 validators at data rate limits from 1000 kbit/s to no limit (i.e., approx 2.2 Mbit/s).

Medium degrees of decentralisation (12–40 validators) show latencies between 6.9 s (12 validators, no limit) and 120.6 s (40 validators, 1000 kbit/s). In general, the latencies rise with increasing degrees of decentralisation and vary in their dependency with the available data rate. A set of 32 validators is limited by a 1000 kbit/s data rate and takes 20.2 s to produce a block, but stabilises at approx. 10 s with data rates of 1500 kbit/s and up. A set of 40 validators, however shows a high dependency on the data rate, lowering from 120.5 s at 1000 kbit/s to 27.8 s at no data rate limit.




Fig. 10. Mean latency of final transactions at maximum transaction throughput at data rates from 50–750 kbit/s and decentralisation degrees from 1–12 validators. The tests show the average values from five repetitions.

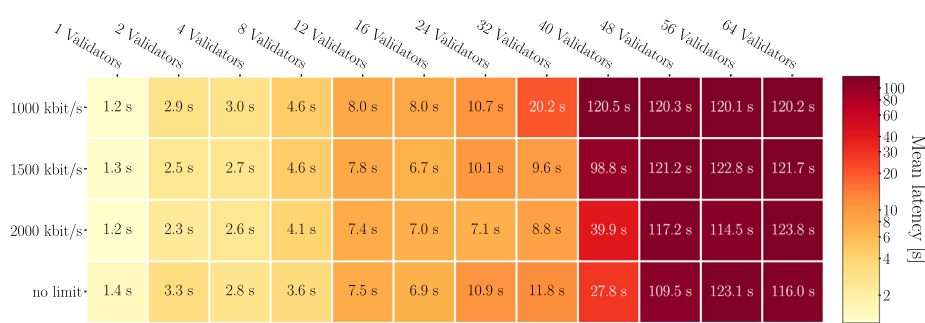


Fig. 11. Mean latency of final transactions at maximum transaction throughput at data rates from 1000 kbit/s to no limit, and decentralisation degrees from 1–64 validators. The results show the average values from five repetitions.

High degrees of decentralisation (40–64 validators) do not achieve latencies under 100 s. This result strengthens the assumptions that the consensus process is computationally too intensive for the Raspberry Pi at high degrees of decentralisation.

6. Discussion

In this section, we discuss how the presented results can be leveraged to determine the infrastructure for blockchain-based microgrids. More specifically, we present a three step procedure to derive meaningful guidelines in regard to the operation of a local energy market, its degree of decentralisation, and the choice of infrastructure from the results presented in Section 5. Furthermore, we reflect this approach on the basis of the case study and assertions of the Walenstadt microgrid, and discuss applicability and limitations of the presented results.

6.1. Designing blockchain-based microgrid infrastructures

In order to utilise the maximum transaction throughput values and relate them to the participant numbers and clearing intervals of the local energy market an approach has been developed that includes the following three steps:

1. Calculate the required transactions per second for a given number of participants and a real-time market clearing interval
2. Decide on the degree of decentralisation of the system (i.e., how many validators should run the blockchain platform and market application)
3. Deduce the necessary minimum data rate requirements and choose appropriate communication infrastructure

Table 3

Required transaction rates, in transactions per second, of the utilised market application for different clearing interval times, and participant numbers.

Clearing interval [s]	Closing time [s]	Participants				
		75	150	300	600	900
60	6	1.4 tps	2.8 tps	5.6 tps	11.1 tps	22.2 tps
300	30	0.3 tps	0.6 tps	1.1 tps	2.2 tps	4.4 tps
600	60	0.1 tps	0.3 tps	0.6 tps	1.1 tps	2.2 tps
900	90	0.1 tps	0.2 tps	0.4 tps	0.7 tps	1.5 tps

1. *Required transactions per second.* In order to calculate required transactions per second of the LEM, we consider the clearing interval and the number of participants. In its current configuration, mentioned in Section 3.2, the deployed market application clears at an interval of 15 min (= 900 s), with a closing time of 90 s (10 %). Therefore, the actual bidding phase is 900 s – 90 s = 810 s long. Within this time, each participant sends a transaction containing its bid to the system, which is required for processing. The required rate of transactions, $r_{transactions}$ in transaction per second, can then be calculated with:

$$\frac{n_{participants}}{(t_{clearing} - t_{closing})} = r_{transactions} \quad (1)$$

where $n_{participants}$ is the number of participants in the microgrid, $t_{clearing}$ is the clearing interval in seconds in which the market is cleared, and $t_{closing}$ is the closing time in seconds after which bids are no longer accepted to the market. Common clearing intervals for LEMs include 600 s, 300 s, and 60 s [37]. Table 3 shows the values of Eq. (1) for clearing interval times from 60–900 s. The values for participant numbers range from the actual size of the Walenstadt microgrid field-test (75) up to 900, in order to show the upper scalability limits.

2. *Degree of decentralisation.* The degree of decentralisation is a design choice for the operation of a decentralised system. In a centralised

system, a single entity operates the application logic (the local energy market). For blockchain-based systems, the degree of decentralisation is determined by the amount of validators, which participate in the consensus process of the underlying blockchain. In a fully decentralised configuration, every participant acts as a validator and contributes to the secure operation of the system. As fully decentralised systems have been demonstrated to require intensive computation and communication, a trade-off decision between security and infrastructural support must be made.

3. Infrastructure choice. Based on the requirements of the local energy market and the number of validators of the underlying blockchain platform, the provided results can be used to derive the minimal data rate requirements, as well as the resulting latency of the system. The infrastructure has to provide the minimal requirements, and can be then weighed off by further influence factors like availability or installation cost.

Walenstadt microgrid. Based on the case study of the Walenstadt microgrid, we can apply the following parameters using the described process and comment on the feasibility of the decentralised local energy market within the Walenstadt microgrid system:

A participant size of 75 devices, and a 900 s clearing interval, results in a minimum requirement of >0.1 tps, which must be processed by the blockchain platform running the market application. As mentioned in Section 3.3, the blockchain platform and market application is run by assigning each producing participant as a system validator. In case of the Walenstadt microgrid, the number of devices representing prosumer PV systems result in 27 validators (please refer to Section 3.1 for more information). The measured results in Fig. 9, show that a system sized between 24–32 validators can sustain over 8 times the required transaction throughput of 0.1 tps at a data rate limit of 1500 kbit/s. This data rate amounts to a sixfold increase over the required data rate of 250 kbit/s determined for a LEM of many of the smart grid applications from the papers described in Section 2.2 (i.e., Ref. [28,33]). When connected with the maximum available data rate (no limit) of the coaxial cable network, this configuration provides a maximum of transaction throughput between 1.8 tps (32 validators, no limit) and 2.2 (24 validators, no limit). With a security factor over 18 times the required transaction throughput, the system design is well within its capabilities to support the required throughput. According to the measured latencies (available in Fig. 11) the local energy market operates with a block time of approximately 12 s.

However, when comparing the case of a 15 min clearing interval real-time LEM with increasing participant numbers and a constant number of validators, the Walenstadt microgrid configuration (27 validators, no limit data rate) could support up to 600 participants (requires 0.7 tps) with a security factor of over 2x. When lowering the clearing interval for a market with a higher temporal resolution, the configuration can maintain a 60 s clearing interval for 75 participants (requires 1.4 tps), but offers security factors of only 1.2x to 1.5x.

6.2. Limitations and future work

The presented results and the provided process to obtain matching communication infrastructure constitutes a mechanism for the evaluation of blockchain-based local energy market systems. The approach is based on an application-specific blockchain platform and field-test specific hardware. The applied methodology provides a process to benchmark blockchain-based smart grid applications, and shows upper limits of processing throughput and latency, which occur due to the degrees of decentralisation (number of validators) and infrastructural limitations (maximum available data rate, maximum tolerable latency).

Application and implementation. The energy market application requires a single transaction per participant and clearing interval. This property is specific to the presented real-time LEM, but can be adapted to any specific market or control application based on messages or transactions.

The implementation of the introduced local energy market is done in JavaScript executed within the NodeJS runtime environment. The computational limits imposed by the Raspberry Pi hardware platform are bound to the efficiency and speed of the implementation. These limits can be countered by moving the implementation of the application from the interpreted language JavaScript to a compiled language, such as Go. However, the shown data rate limitations are valid, regardless of the implementation language.

Consensus mechanism. The Tendermint consensus mechanism used to coordinate the decentralisation of the platform is not explicitly built for the use in IoT devices and has a high overhead when running with as many validators as was demonstrated in the conducted tests. The tests, however, provide an orientation of expected throughput at different degrees of decentralisation and data rates.

Extensibility. The shown results are specific to the introduced market application. Additional smart grid applications, such as Demand Side Management (DSM), when based on a fully consensus afflicted process, require higher amounts of information from the participants. This information is delivered to the control application in form of transactions, which should be considered when applying the provided results to alternative applications and scenarios.

Smart metre infrastructure. The hardware capabilities of smart-metres on the end-customer level are nowhere near those of the utilised Raspberry Pi SBC hardware in the Walenstadt microgrid field test. Standard smart-metres enable remote readings by connected devices, but do not offer an environment to run custom software like the system presented in this paper. Utility companies could further push the development of smart-metres in order to enable customisability.

Future work. In real-world operation, each of the participant devices is subject to downtime. The conducted tests do not encompass the system behaviour in terms of failures of agents and validators. In the future, the conducted tests can be extended to focus not only on maximum throughput, but also on the reliability of the system and stability during intermittent communication, blackouts, and behaviour during cold starts. Future research initiatives could focus on which consensus mechanisms favour consistency over availability and which are the best choice for smart grid applications.

In order to include smart grid services with high transaction throughput and low latency requirements, such as Demand Response (DR) or Distribution Automation (DA), off-chain scaling solutions (payment channels) could be subject to future testing. These solutions provide excellent scaling behaviour by handling certain processes without consensus, while using blockchain-based coordination for selected actions, such as settlements and milestones [44].

7. Conclusion

As an emergent technology, blockchain has drawn considerable attention from the utility industry, associated service providers, and academia. This novel technology has let adopters explore novel use cases and concepts, which are believed to enable novel business solutions like peer-to-peer (P2P) trading. Although numerous articles have worked out conceptual foundations and benefit/risk analyses of blockchain technology, several pilot-projects have been conducted in the past years with a focus on P2P trading. However, only a small number of these pilot tests were conducted in an academic setting and led to subsequent academic publications with presented insights that go beyond conceptual studies, simulations, and requirements engineering.

In this article, we present the concept, design and implementation details of the Walenstadt microgrid (a.k.a. “Quartierstrom”), a real-world field test with 37 participating households (27 of which are prosumers). A total of 75 smart-metres with integrated Raspberry Pi Single Board Computers (SBCs) were deployed to measure demand, photovoltaic power output, and battery loads. The SBCs are used to host an application specific consortium blockchain which implements a double auction market place, in which prosumers and consumers can adjust their price preferences for bids and asks. The blockchain allows the nodes to reach agreement over the computed market prices and involved transactions for each trading period. The Walenstadt microgrid is perfectly suited to study different design configurations regarding number of validators and available network infrastructure. The field setup is utilised to determine how the number of validators and the maximum data rate of the underlying communication infrastructure affect the transaction throughput and system latency. Conclusions regarding scalability based on achieved throughput under different design configurations are presented as follows.

The tests include artificially limited data rates from 50–16 000 kbit/s as well as a no limit. The open data rate limit of the available coaxial cable connection was measured to be approximately 2.2 Mbit/s. The dependent variables maximum transaction throughput in transactions per second (tps) and latency in seconds were measured for configurations of 1–64 validators. We tested the feasibility of the project setup with its communication requirements and maximum throughput limits. Although the tests show that transaction throughput based on the hardware used in the field test are below 10 tps, and therefore in the range of the maximum throughput achieved by public blockchains, it is still feasible to run a blockchain-based local energy market with them. We found that the operation cannot be guaranteed if the maximum data rate of the communication infrastructure is below 1000 kbit/s, which rules out communication technologies such as Wireless Personal Area Network (WPAN), narrowband Power Line Communication (PLC), and Narrowband Internet of Things (NB-IoT). More generally, the number of validators increased the degree of decentralisation and fault-tolerance, but decreased the throughput due to communication overhead between nodes. The tests demonstrated that having more than 40 validator nodes will throttle throughput to the extent that the operation of the blockchain became infeasible. We found that the system implemented in Walenstadt, with its requirements for a high degree of decentralisation and its underlying infrastructure, is able to provide 17 times the required throughput.

The decentralisation of energy resources requires novel approaches to control the grid infrastructure. Microgrid structures are expected to be a basic feature of future distribution networks and the smart grid, in order to take full advantage of DERs. Local Energy Markets (LEMs) provide mechanisms to provide financial settlement between micro-generation owners, storage owners, consumers, and embed microgrids into existing market structures. Blockchain technology has shown a promising rise in academic research with its application of trusted, decentralised execution of peer-to-peer (P2P) markets, as well as control mechanisms.

This paper represents the next steps to making these systems a reality. With these systems, local prosumers are able to achieve a higher remuneration rate for their generated electricity and consumers have the benefit of using locally generated electricity and reduce dependence on the centralised grid. The technology enabling this collaboration is based on computationally highly capable smart-metres, which are connected in a Virtual Private Network, and the software that runs on them. We introduced the deployed mechanisms and data structures behind the local energy market application. Furthermore, we introduced the blockchain platform, which enables the market application to run in a decentralised setting and in a consensus which is enabled by a trusted share of the microgrid’s participants. In order to provide a method of analysis for the overall system performance and requirements on the underlying communication infrastructure, we present a methodology

to perform throughput and latency tests on the platform. The tests are based on blockchain performance tests and provide guidelines to apply the tests to other blockchain-based local energy market platforms.

The Walenstadt microgrid pilot project shows that these systems can work successfully today, and with improvement, they can change the structure of our electricity grids from a centralised form of production, to a decentralised form that benefits prosumers and integrates more renewable energy to decarbonise our electricity grid.

CRedit authorship contribution statement

Arne Meeuw: Conceptualization, Methodology, Software, Writing - original draft. **Sandro Schopfer:** Validation, Visualization, Writing - review & editing. **Anselma Wörner:** Writing - review & editing. **Verena Tiefenbeck:** Writing - review & editing. **Liliane Ableitner:** Writing - review & editing. **Elgar Fleisch:** Supervision. **Felix Wortmann:** Supervision, Writing - review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

Funding

This project is being supported by the Swiss Federal Office of Energy (SFOE) within the framework of its pilot, demonstration and flagships programme.

References

- [1] US Energy Information Administration, More than Half of Small-Scale Photovoltaic Generation Comes from Residential Rooftops, Tech. rep., 2017, URL <https://www.eia.gov/todayinenergy/detail.php?id=31452>.
- [2] Fraunhofer ISE, H. Wirth, Aktuelle Fakten Zur Photovoltaik in Deutschland (Amended on 21st of 2018), Tech. rep., 2018, URL <https://www.ise.fraunhofer.de/content/dam/ise/de/documents/publications/studies/aktuelle-fakten-zur-photovoltaik-in-deutschland.pdf>.
- [3] S. Schopfer, V. Tiefenbeck, T. Staake, Economic assessment of photovoltaic battery systems based on household load profiles, Appl. Energy 223 (2018) 229–248, <http://dx.doi.org/10.1016/j.apenergy.2018.03.185>, URL <http://www.sciencedirect.com/science/article/pii/S0306261918305026>.
- [4] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, A. Peacock, Blockchain technology in the energy sector: A systematic review of challenges and opportunities, Renew. Sustain. Energy Rev. 100 (2019) 143–174, <http://dx.doi.org/10.1016/j.rser.2018.10.014>, URL <http://www.sciencedirect.com/science/article/pii/S1364032118307184>.
- [5] O. Utility, Whitepaper: A glimpse into the future of Britain’s energy economy, 2016, URL piclo.energy.
- [6] Vandebroun - Duurzame energie van Nederlandse bodem. URL <https://vandebron.nl/>.
- [7] C. Wang, J. Yan, C. Marnay, N. Djilali, E. Dahlquist, J. Wu, H. Jia, Distributed energy and microgrids (DEM), Appl. Energy 210 (2018) 685–689, <http://dx.doi.org/10.1016/j.apenergy.2017.11.059>, URL <http://linkinghub.elsevier.com/retrieve/pii/S0306261917316550>.
- [8] C. Shen, F. Pena-Mora, Blockchain for cities - A systematic literature review, IEEE Access (2018) 1, <http://dx.doi.org/10.1109/ACCESS.2018.2880744>.
- [9] T.T.A. Dinh, J. Wang, G. Chen, R. Liu, B.C. Ooi, K.-L. Tan, BLOCK-BENCH: A Framework for Analyzing Private Blockchains, ACM, 2017, pp. 1085–1100, <http://dx.doi.org/10.1145/3035918.3064033>, URL <http://dl.acm.org/citation.cfm?id=3035918.3064033>.
- [10] D. Drescher, A NON-TECHNICAL INTRODUCTION IN 25 STEPS 250.
- [11] K. Christidis, M. Devetsikiotis, Blockchains and smart contracts for the internet of things, IEEE Access 4 (2016) 2292–2303, <http://dx.doi.org/10.1109/ACCESS.2016.2566339>.
- [12] S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System. URL <http://www.bitcoin.org/bitcoin.pdf>.
- [13] others, A next-generation smart contract and decentralized application platform, white paper.

- [14] L. Consensus, Grid+ White Paper, Tech. rep., 2018, URL <https://gridplus.io/assets/Gridwhitepaper.pdf>.
- [15] P.L. Power Ledger, Power Ledger White Paper, Tech. rep., 2018, URL <https://powerledger.io/media/Power-Ledger-Whitepaper-v8.pdf>.
- [16] L. Energy, Brooklyn Microgrid Overview, 2019, URL <http://brooklynmicrogrid.com/>.
- [17] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, A. Peacock, Blockchain technology in the energy sector: A systematic review of challenges and opportunities, *Renew. Sustain. Energy Rev.* 100 (2019) 143–174, <http://dx.doi.org/10.1016/j.rser.2018.10.014>, URL <http://www.sciencedirect.com/science/article/pii/S1364032118307184>.
- [18] E. Mengelkamp, J. Gärtner, K. Rock, S. Kessler, L. Orsini, C. Weinhardt, Designing microgrid energy markets, *Appl. Energy* 210 (2018) 870–880, <http://dx.doi.org/10.1016/j.apenergy.2017.06.054>, URL <http://linkinghub.elsevier.com/retrieve/pii/S030626191730805X>.
- [19] J.J. Sikorski, J. Houghton, M. Kraft, Blockchain technology in the chemical industry: Machine-to-machine electricity market, *Appl. Energy* 195 (2017) 234–246, <http://dx.doi.org/10.1016/j.apenergy.2017.03.039>, URL <http://www.sciencedirect.com/science/article/pii/S0306261917302672>.
- [20] AspenTech, Aspen plus, 2019, URL <https://www.aspentech.com/products/engineering/aspen-plus/>.
- [21] A. Goranovic, M. Meisel, L. Fotiadis, S. Wilker, A. Treytl, T. Sauter, Blockchain Applications in Microgrids an Overview of Current Projects and Concepts, *IEEE*, 2017, pp. 6153–6158, <http://dx.doi.org/10.1109/IECON.2017.8217069>, URL <http://ieeexplore.ieee.org/document/8217069/>.
- [22] S. Cheng, B. Zeng, Y.Z. Huang, Research on application model of blockchain technology in distributed electricity market, in: *IOP Conference Series: Earth and Environmental Science*, Vol. 93, 2017, 012065, <http://dx.doi.org/10.1088/1755-1315/93/1/012065>.
- [23] J. Wang, Q. Wang, N. Zhou, Y. Chi, A novel electricity transaction mode of microgrids based on blockchain and continuous double auction, *Energies* 10 (12) (2017) 1971, <http://dx.doi.org/10.3390/en10121971>, URL <https://www.mdpi.com/1996-1073/10/12/1971>.
- [24] C. Pop, T. Cioara, M. Antal, I. Anghel, I. Salomie, M. Bertoncini, C. Pop, T. Cioara, M. Antal, I. Anghel, I. Salomie, M. Bertoncini, Blockchain based decentralized management of demand response programs in smart energy grids, *Sensors* 18 (1) (2018) 162, <http://dx.doi.org/10.3390/s18010162>, URL <https://www.mdpi.com/1424-8220/18/1/162>.
- [25] F. Lombardi, L. Aniello, S.D. Angelis, A. Margheri, V. Sassone, A blockchain-based infrastructure for reliable and cost-effective IoT-aided smart grids, in: *Living in the Internet of Things: Cybersecurity of the IoT - 2018*, 2018, pp. 1–6, <http://dx.doi.org/10.1049/cp.2018.0042>.
- [26] C. Greer, D.A. Wollman, D.E. Prochaska, P.A. Boynton, J.A. Mazer, C.T. Nguyen, G.J. FitzPatrick, T.L. Nelson, G.H. Koepke, A.R.H. Jr, V.Y. Pillitteri, T.L. Brewer, N.T. Gormie, D.H. Su, A.C. Eustis, D.G. Holmberg, S.T. Bushby, NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 3.0, Special Publication (NIST SP) - 1108r3, doi: <http://dx.doi.org/10.1002/http://dx.doi.org/10.6028/NIST.SP.1108r3>, URL [https://www.nist.gov/publications/nist-framework-and-roadmap-smart-grid-interoperability-st\(ands\)-release-30](https://www.nist.gov/publications/nist-framework-and-roadmap-smart-grid-interoperability-st(ands)-release-30).
- [27] Y. Liang, F. Liu, C. Wang, S. Mei, Distributed demand-side energy management scheme in residential smart grids: An ordinal state-based potential game approach, *Appl. Energy* 206 (2017) 991–1008, <http://dx.doi.org/10.1016/j.apenergy.2017.08.123>, URL <http://www.sciencedirect.com/science/article/pii/S0306261917311467>.
- [28] E. Ancillotti, R. Bruno, M. Conti, The role of communication systems in smart grids: Architectures, technical solutions and research challenges, *Comput. Commun.* 36 (17) (2013) 1665–1697, <http://dx.doi.org/10.1016/j.comcom.2013.09.004>, URL <http://www.sciencedirect.com/science/article/pii/S0140366413002090>.
- [29] M. Kuzlu, M. Pipattanasorn, S. Rahman, Communication network requirements for major smart grid applications in HAN, NAN and WAN, *Comput. Netw.* 67 (2014) 74–88, <http://dx.doi.org/10.1016/j.comnet.2014.03.029>, URL <http://linkinghub.elsevier.com/retrieve/pii/S1389128614001431>.
- [30] R.H. Khan, J.Y. Khan, A comprehensive review of the application characteristics and traffic requirements of a smart grid communications network, *Comput. Netw.* 57 (3) (2013) 825–845, <http://dx.doi.org/10.1016/j.comnet.2012.11.002>, URL <http://linkinghub.elsevier.com/retrieve/pii/S1389128612003751>.
- [31] U. DoE, Communications Requirements of Smart Grid Technologies, Tech. Rep, US Department of Energy, 2010, pp. 1–69.
- [32] M.F. Zia, E. Elbouchikhi, M. Benbouzid, Microgrids energy management systems: A critical review on methods, solutions, and prospects, *Appl. Energy* 222 (2018) 1033–1055, <http://dx.doi.org/10.1016/j.apenergy.2018.04.103>, URL <http://www.sciencedirect.com/science/article/pii/S0306261918306676>.
- [33] Y. Li, X. Cheng, Y. Cao, D. Wang, L. Yang, Smart choice for the smart grid: Narrowband internet of things (NB-IoT), *IEEE Internet Things J.* 5 (3) (2018) 1505–1515, <http://dx.doi.org/10.1109/JIOT.2017.2781251>.
- [34] W. Giezmann, J. Stokking, The things network, 2019, URL <https://client:3000/>.
- [35] E. Buchman, Tendermint: Byzantine Fault Tolerance in the Age of Blockchains.
- [36] R. Han, V. Gramoli, X. Xu, Evaluating blockchains for IoT, in: 2018 9th IFIP International Conference on New Technologies, Mobility and Security (NTMS), 2018, pp. 1–5, <http://dx.doi.org/10.1109/NTMS.2018.8328736>.
- [37] F. Blom, H. Farahmand, On the scalability of blockchain-supported local energy markets, in: 2018 International Conference on Smart Energy Systems and Technologies (SEST), 2018, pp. 1–6, <http://dx.doi.org/10.1109/SEST.2018.8495882>.
- [38] Enerserve, SmartPi 2.0, 2019, URL <https://shop.enerserve.eu/smartpi/262/smartpi-2.0>.
- [39] L. Lamport, M. Fischer, Byzantine Generals and Transaction Commit Protocols, Tech. rep., Technical Report 62, SRI International, 1982.
- [40] J. Kwon, Tendermint: Consensus Without Mining, 2014, p. 11.
- [41] S. Chivtanasantorn, B. Bhumirata, E. Filiol, Perseus on VoIP: Development and implementation of VoIP platforms, in: 2014 11th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON), 2014, pp. 1–5, <http://dx.doi.org/10.1109/ECTICon.2014.6839737>.
- [42] H. Hong, From cloud computing to fog computing: Unleash the power of edge and end devices, in: 2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), 2017, pp. 331–334, <http://dx.doi.org/10.1109/CloudCom.2017.53>.
- [43] Tendermint, Tendermint | tendermint documentation, 2019, URL <https://tendermint.com/docs/>.
- [44] R. Khalil, A. Gervais, Revive: Rebalancing off-blockchain payment networks, in: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17, ACM, New York, NY, USA, 2017, pp. 439–453, <http://dx.doi.org/10.1145/3133956.3134033>, event-place: Dallas, Texas, USA. URL <http://doi.acm.org/10.1145/3133956.3134033>.

Glossary

Acronyms

AMI: Advanced Metering Infrastructure
 BFT: Byzantine Fault Tolerance
 CPU: Central Processing Unit
 DA: Distribution Automation
 DER: Distributed Energy Resource
 DLT: Distributed Ledger Technology
 DR: Demand Response
 DRM: Demand Response Management
 DSL: Digital Subscriber Line
 DSM: Demand Side Management
 DSO: Distribution System Operator
 HAT: Hardware Attached on Top
 HLF: Hyperledger Fabric
 HTTP: Hypertext Transfer Protocol
 IoT: Internet of Things
 JSON: JavaScript Object Notation
 LAN: Local Area Network
 LEM: Local Energy Market
 LoRaWAN: Long Range Wide Area Network
 LP-WPAN: Low Power Wireless Personal Area Network
 NAN: Neighbourhood Area Network
 NB-IoT: Narrowband Internet of Things
 P2P: peer-to-peer
 PBFT: Practical Byzantine Fault Tolerance
 PLC: Power Line Communication
 PoA: Proof of Authority
 PoS: Proof of Stake
 PoW: Proof of Work
 PV: photovoltaic
 RAM: Random Access Memory
 RPC: Remote Procedure Call
 RTP: real-time pricing
 SBC: Single Board Computer
 TCP: Transport Control Protocol
 TOU: time-of-use
 tps: transactions per second
 TTN: The Things Network
 VPN: Virtual Private Network
 WPAN: Wireless Personal Area Network