



# Research Repository UCD

<b>Title</b>	A Delay-Tolerant Payment Scheme Based on the Ethereum Blockchain
<b>Authors(s)</b>	Hu, Yining, Manzoor, Ahsan, Ekparinya, Parinya, Liyanage, Madhusanka, et al.
<b>Publication date</b>	2019-03-06
<b>Publication information</b>	Hu, Yining, Ahsan Manzoor, Parinya Ekparinya, Madhusanka Liyanage, and et al. "A Delay-Tolerant Payment Scheme Based on the Ethereum Blockchain." IEEE, March 6, 2019. <a href="https://doi.org/10.1109/ACCESS.2019.2903271">https://doi.org/10.1109/ACCESS.2019.2903271</a> .
<b>Publisher</b>	IEEE
<b>Item record/more information</b>	<a href="http://hdl.handle.net/10197/11692">http://hdl.handle.net/10197/11692</a>
<b>Publisher's statement</b>	© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
<b>Publisher's version (DOI)</b>	10.1109/ACCESS.2019.2903271

Downloaded 2025-12-04 22:46:23

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd\_oa)



© Some rights reserved. For more information

# A Delay-Tolerant Payment Scheme Based on the Ethereum Blockchain

Yining Hu,<sup>\*†</sup> Ahsan Manzoor,<sup>‡</sup> Parinya Ekparinya,<sup>§</sup> Madhusanka Liyanage,<sup>‡</sup>  
Kanchana Thilakaratna,<sup>§</sup> Guillaume Jourjon,<sup>†</sup> Aruna Seneviratne,<sup>\*†</sup> and Mika E Ylianttila<sup>‡</sup>

<sup>\*</sup>University of New South Wales, Australia, <sup>†</sup>Data61-CSIRO, Australia,

<sup>‡</sup>University of Oulu, Finland, <sup>§</sup>University of Sydney, Australia

Email: <sup>\*</sup>firstname.lastname@data61.csiro.au, <sup>†</sup>firstname.lastname@oulu.fi,

<sup>§</sup>pekp6601@uni.sydney.edu.au, <sup>§</sup>kanchana.thilakaratna@sydney.edu.au

**Abstract**—Banking as an essential service can be hard to access in remote, rural regions where the network connectivity is intermittent. Although micro-banking has been made possible by SMS or USSD messages in some places, their security flaws and session-based nature prevent them from a wider adoption. Global level cryptocurrencies enable low-cost, secure and pervasive money transferring among distributed peers, but are still limited in their ability to reach more people in remote communities.

We proposed to take advantage of the delay-tolerant nature of blockchains to deliver banking services to remote communities that only connect to the broader Internet intermittently. Using a base station that offers connectivity within the local area, regular transaction processing is solely handled by blockchain miners. The bank only joins to process currency exchange requests, reward miners and track user balances when the connection is available. By distributing the verification and storage tasks among peers, our system design saves on the overall deployment and operational costs without sacrificing the reliability and trustworthiness. Through theoretical and empirical analysis, we provided insights to system design, tested its robustness against network disturbances, and demonstrated the feasibility of implementation on off-the-shelf computers and mobile devices.

## I. INTRODUCTION

Business operations in rural areas often face challenges associated with access, due to the lack of infrastructures. Communication infrastructure is one of the most crucial challenges, as an increasing number of services today rely on continuous network connectivity. However, connectivity in many of those difficult-to-access areas, even when available, is often intermittent. As a consequence, most of the pervasive services that are taken for granted cannot operate in these connectivity-restricted environments. This has led to the development of techniques, such as ad-hoc networks, and applications that can utilize the intermittent connection to provide some level of service [29], [25] in the past. More recently, these regions are beginning to be served by community-run small-cell base stations for local connectivity, for example, Nokia Kuha [1] base stations that enable connectivity to the broader internet intermittently, but continuous interconnectivity for devices located within its coverage.

One sector that has adopted different technologies available in these connectivity-restricted environments to deliver valuable services, is banking. For example, primitive banking services have been made possible via Short Message Service (SMS) or Unstructured Supplementary Service Data (USSD) of cellular networks. This is exemplified by the BAAC in Thailand [11] and M-Pesa in Kenya and Tanzania [21]. The emergence of having interconnectivity within a given area, e.g., a village, with intermittent access to the main operator

network, as well as decentralised services such as cryptocurrencies that do not require centralized control [19], [20], has opened up possibilities for delivering improved banking services to connectivity-restricted areas.

Indeed, cryptocurrencies have received a significant attention in recent years and are regarded by many as a potential revolution of the banking industry [19], [20]. The technology behind, the blockchain, is secured by hard-coded software programs and enables peer democracy for transaction settlement. However, like most other pervasive services of today, cryptocurrencies require continuous network connectivity to constantly exchange large volumes of data. For instance, a typical Bitcoin client uses around 200 MB of data per hour [2], while the Ethereum blockchain size has even surpassed the Bitcoin's in June 2017.<sup>1</sup> Therefore, to take advantage of the village-wide interconnectivity and the benefits of decentralisation, it is necessary to leverage the delay-tolerant characteristic of the blockchain and extend it to operate in intermittently connected environments.

This paper shows how this can be achieved by considering a scenario where a community operates small-cell base stations for local connectivity with intermittent connection to the broader Internet, and a financial institution, i.e., a bank, that supplies, pays for and monitors equipment used in providing the service and makes the following contributions:

- 1) The design of a low-cost, accessible, reliable and secure payment scheme based on the Ethereum blockchain.
- 2) Mathematical models of the proposed system and an evaluation of the system performance under several operating conditions that confirms its feasibility.
- 3) Demonstrate the practical viability of proposed system through a prototype implementation using off-the-shelf laptops and mobile devices.
- 4) The first ever use of smart contracts for admission control, management of user accounts, mining rewards, and token creation.

The remainder of the paper is organized as follows: Section II introduces technical details of blockchain. Section III describes the system architecture and Section IV provides models for the system design. Section V presents the validation of our models and evaluation over local blockchain system. Section VI demonstrates a prototype implementation with multiple test results. Section VII discusses the related work, Section VIII highlights insights on future directions, and Section IX finally concludes the paper.

<sup>1</sup><http://www.altcointoday.com/ethereums-blockchain-size-surpasses-bitcoins-by-40/>

## II. BLOCKCHAIN BASICS

We base our system design on the Ethereum platform [31] that incorporates a Turing-complete scripting language for user-defined programs with a short block-generation time. We here explain key terminology and describe the Ethereum mechanisms.

### A. Nodes and Their Connectivity

A blockchain consists of a peer-to-peer (P2P) communication overlay network. In the case of Ethereum, each network node continuously attempts to connect to other nodes until they have peers. By default, Ethereum nodes use a gossip protocol to find out about other nodes.

Once the overlay network is created, nodes on a blockchain may act as full nodes, miners or light nodes. All nodes contribute to the network connectivity and information (i.e., transactions and blocks) propagation. Miners and full nodes verify all transactions based on the signatures on them and the state changes they result in, and each keeps a complete transaction record. Light nodes operate in the Simplified Payment Verification (SPV) mode to only download block headers and verify and store transactions that are related to them. In addition, miners settle transactions through a process called mining as explained in Part II-C.

### B. Transactions, Blocks and Smart Contracts

The state of a blockchain is represented by peer interactions or transactions. For simplicity and efficiency, transactions are grouped together to be settled and immutably recorded in blocks. A user needs an externally controlled account (EOA) to send and receive transactions. An EOA has an ether balance and is controlled by the user's private key. Ethereum also uses smart contracts to form agreements among different entities. Smart contracts are identified by contract accounts, which are similar to EOAs but are associated with code that can be triggered by transactions or messages (calls) received from other contracts. Transactions can be sent between two EOAs, two contract accounts, or an EOA and a contract account. In the public Ethereum network, miners also collect transaction fees to make profit [31]. Transaction priority, i.e. how likely a transaction is to be picked up by miners, depends mostly on its value, age and the transaction fee attached. State of a blockchain is computed after every block [4], and smart contracts are executed by all nodes when synchronization happens (cf. Part II-C).

### C. The Blockchain

Nakamoto in his original paper proposed a proof-of-work (PoW) scheme to confirm transactions and select representation in majority decision making [22]. PoW is also adopted by the current version of Ethereum as described below.

1) *PoW Mining and Difficulty Control*: PoW mining is essentially the process of repeatedly calculating a hash value with an incrementing nonce until the hash is smaller than a target. Hashrate is the speed at which a miner computes hashes. Miners compete against each other in solving PoW puzzles and broadcast blocks to the rest of the network upon block creation for validation and synchronization. In addition to mining, miners also listen for new transactions and blocks discovered by others to prepare the next block.

Difficulty is a measure of how difficult it is to solve a PoW puzzle. Difficulty adjustment over PoW puzzles leads to a stable block creation rate. In Ethereum the calculation is based on the block number, timestamp of the current block, and timestamp & difficulty of its parent block. For Ethereum Homestead,<sup>2</sup> the adjustment is always a multiple of  $\text{parent\_diff} // 2048$ , with parameter  $a \in [-99, 1]$ . We denote  $\text{block\_timestamp} - \text{parent\_timestamp}$  as  $\delta t$ , and summarize the adjustments in Table I. Difficulty adjustment is insignificant under a high network hashrate.

Table I: PoW difficulty adjustments.

Timestamp difference	Adjustment
$\delta t < 10s$	$a = 1$
$10s \leq \delta t < 20s$	Difficulty unchanged, $a = 0$
$20s \leq \delta t < 1000s$	Adjust downwards proportional to $\delta t$ , $a \in [-99, -1]$
$\delta t \geq 1000s$	$a = -99$

2) *Consensus and Chain Growth*: Transmission delays may lead to stale blocks, or forks [9] as multiple blocks could be created at the same time, and received by different nodes across the network. A block is attached to the chain once created, with its own block header pointing to the previous block header, and all the way back to the genesis block. All peers work out these inconsistencies by selecting the longest chain. Ethereum determines the longest chain based on the total difficulty of all the blocks.<sup>3</sup>

3) *Node Synchronization*: Network delays and node churns are common in blockchain systems. If one node restarts after going off-chain, it enquires its peers to obtain the latest version of the ledger. States stored in a particular block can only be accessed if the node synchronization has reached it.

### D. Coin Supply

To compensate the resources consumed in mining, a fixed amount of new tokens, or mining reward, is generated upon the creation of blocks and paid to the winners. This rewarding scheme is necessary as it makes the blockchain more secure [5], but the continuous creation of coins may result in inflation.

## III. SYSTEM ARCHITECTURE

The primary objective of the proposed system is to enable cash-less payment in a remote community that has constant connectivity within the local area but only intermittent connection to the broader Internet. Namely, the community connects to a bank's central network intermittently. A network operator deploys a base station, similar to Nokia Kuha [1] that forms a 4G LTE cellular network in the local region. The backhaul network could be a low-bandwidth satellite or a long-range microwave connection that does not provide robust service guarantees. Transaction processing is enabled using a private Ethereum blockchain [31] deployed by the bank. We assume some villagers would volunteer to participate in mining and be chosen by the bank. We do not use the default ether generation scheme to avoid inflation in the local economy, but instead, create our own digital currency, or Token,<sup>4</sup> via

<sup>2</sup><http://ethdocs.org/en/latest/introduction/the-homestead-release.html>

<sup>3</sup>Ethereum does not implement the full Greedy Heaviest-Observed Sub-Tree (GHOST) protocol proposed by Sompolinsky and Zohar [30], instead, it uses a longest-chain rule with rewards for stale blocks [12].

<sup>4</sup><https://www.ethereum.org/token>

a smart contract, for use within the community. The use of Tokens automatically imposes restrictions on users entering the system and makes it easier for the bank to manage user accounts.

Figure 1a graphically illustrates our system design. The bank authorizes a set of selected villagers to act as miners and compensates them by assigning a certain number of Tokens for each valid block they create. Regular users can join the system by running a light node that is deployed as a customized mobile app. All villagers, especially vendors, can host a full node as a payment gateway to keep a complete transaction record if they wish so. The detailed design of the mining network, with the appropriate number of miners and the minimal connectivity requirements for the system to provide the intended service, is described in Section IV. The distinguishing feature of the proposed system is that it enables villagers with Tokens to perform *regular transactions* ( $tx_{P,P}$  in Figure 1b) among themselves with all *regular transactions* treated equally, processed and authorised by the blockchain miners irrespective of the bank's presence.

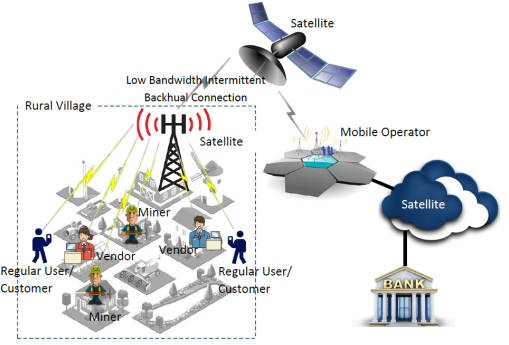
The bank itself hosts a passive full node to avoid creating forks during its disconnection periods. It deploys a smart contract, i.e. the *user balance contract*, to record user balances in both fiat currency and digital currency, as well as distributing mining rewards. When the connection between the bank and the village is established, the bank synchronizes with other blockchain nodes, updates user balances and processes currency exchange requests. These *currency exchange transactions* ( $tx_{B,P}$  in Figure 1b) between the bank and villagers are treated with a higher priority by the blockchain network as they are only committed when the connection is available. The bank also takes appropriate actions in case of malicious behaviors such as suspicious transactions and misbehaving miners. As a result, we have developed the first ever smart contract responsible for the admission control, account management, rewarding miners, and token creation.

Figure 1b illustrates transaction flows in the system. In order to send a transaction, Regular User 1 first requests to load her digital account with the money in her fiat account.<sup>5</sup> The bank then issues transaction  $tx_{B,P}$ , which is picked up by Miner 1 and included in Block #i. Once  $tx_{B,P}$  is confirmed, the code inside the *user balance contract* is triggered and User 1's fiat account and digital account are updated accordingly as shown in Step ①. Next, User 1 sends Vendor 1 a transaction  $tx_{P,P}$  to obtain the service, which is collected by Miner 2 and confirmed in Block #i+1. During the bank's next synchronization, once it reaches Block #i+1, the smart contract code is triggered again and the digital accounts of both User 1 and Vendor 1 are updated as shown in Step ②. The bank also keeps two accounts of itself, just like other blockchain nodes. If an outside party without Tokens wants to make transactions with people in the village, it first submits a request to the bank node in fiat currency. The bank will then convert the payment to the digital currency locally and redirect it to the target receiver when the connection is available.

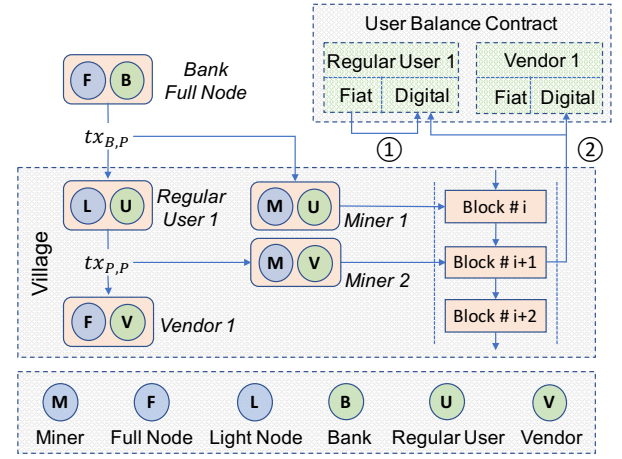
#### IV. SYSTEM MODELLING AND DESIGN

We first probabilistically model transaction processing of the local blockchain system as it is a major concern of users.

<sup>5</sup>A fiat account records the amount of fiat currency a user deposits to the bank.



(a) System architecture.



(b) Transaction Flows.

Figure 1: Overview of the proposed system.

We then model the overall deployment and operational costs from the bank's point of view and use the models to design system properties to enable an efficient and reliable local area payment scheme under the given network constraints.

##### A. Assumptions

We make the following assumptions without loss of generality:

- 1) All individual mining nodes have equal and stable computational power as mining equipment are supplied by the bank to authorised villagers.
- 2) Block size is sufficient to include all transactions for immediate processing as the transaction rates within a village will be considerably lower when compared to a public Ethereum network.
- 3) No transaction fees and no rewards for mining stale blocks as rewarding is controlled by the bank using Tokens.
- 4) Network bandwidth available within the village is sufficient for all blockchain related data traffic.

### B. Modelling Transaction Processing

Regular transaction arrival  $r_{t_i}$ , where  $t_i$  is the transaction initiation time, follows a Poisson distribution as observed in [17]. We denote the average arrival rate as  $\lambda_t$  transactions per second (tps). Block generation time  $T$  has been shown to be exponentially distributed with probability density function  $G(T) = \lambda e^{-\lambda T}$ , where  $E[T] = \frac{1}{\lambda}$  [9]. In Ethereum,  $E[T] \approx 12s$ . Since we assume block size to be large enough (cf. assumption 2), there are no pending transactions as all transactions arriving in the current mining session are included in the next block. Therefore, transaction processing time  $t_p = T - t_i$ , where  $t_i, t_p \in (0, T]$ . We further analyze the distribution of block time and transaction processing time under different arrival rates in Section V-B. Transaction throughput can also be calculated as  $\frac{\sum_{\forall t_i \leq T} r_{t_i}}{T}$ , with transaction size being  $s_t$  bits, blocks need to accommodate  $s_b = \lambda_t s_t E[T]$  bits of transactions on average.

Then, there are *money exchange transactions* with the bank when it is connected. We model money exchange transaction processing similar to that of regular transactions with average arrival rate  $\lambda_e$  and size  $s_e$ .

### C. Modelling the System Cost

We derive cost models from the bank's point of view for capital equipment and operational costs based on metrics in [8]. We omit transaction validation and storage costs in our calculations as they are only a small fraction of mining equipment cost and mining itself [8].

1) *Equipment Cost*: Mining equipment may come in different forms. We use  $d_m$  to denote cost of a mining equipment and  $l_m$  for total number of miners with a life time of  $x_m$  (in years). Then, the averaged equipment cost per year is  $C_E = l_m d_m / x_m$ .

2) *Rewarding Cost*: Although the block reward is an incentive for miners, it is an extra cost for the bank. We use  $R$  to denote the block reward, meaning that bank has to pay  $R$  Tokens for each valid block. We calculate the rewarding cost as  $C_R = nR$ , where  $n$  represents the number of blocks.

3) *Network Resources*: Since we assume ideal network conditions (cf. assumption 4) locally, we only model the network resource usage of the backhaul network. When connected, the bank synchronizes past transactions and processes money exchange requests. We define one service period as  $T_B = T_C + T_U$ , in which  $T_C$  and  $T_U$  stand for the cumulative total duration of connected and disconnected sessions respectively during the service period. We denote the backhaul network bandwidth requirement as  $BW$  and cost as  $C_{BW}$  in Token per bit. Then, the network connectivity cost would be  $C_N = C_{BW} T_C BW$  during one service period.

4) *System Cost*: To sum up, the overall system cost bank has to afford is the sum of all the aforementioned costs. Since cost components are defined with different time units, we calculate the overall cost within time period  $T_0$ , which equals  $x_y$  years,  $x_b$  blocks and  $x_s$  service periods as:

$$\begin{aligned} C_{All} &= C_E x_y + C_R + C_N x_s \\ &= l_m \frac{d_m}{x_m} x_y + R x_b + C_{BW} T_C BW x_s. \end{aligned} \quad (1)$$

### D. Mining Network Design

We now find ranges of multiple mining network parameters including number of miners and their connectivity to ensure reliability and security.

1) *Miner Outages*: Miners are incentivized to work, but they may join or leave the network spontaneously (churn) [14]. We denote total number of online miners as  $l_{on}$ ,  $l_{on} \leq l_m$ . Probability of one miner going offline in each time slot is  $p_1, p_2, p_3, \dots, p_{l_m}$ . Hence the probability that  $X = k$  miners are offline in the same time slot can be represented by a Poisson Binomial Distribution:

$$P_r(X = k) = \sum_{A \in F_k} \prod_{i \in A} p_i \prod_{j \in A^c} (1 - p_j), \quad (2)$$

where Set  $F_k$  contains all subsets of  $k$  integers that can be selected from  $\{1, 2, 3, \dots, l_m\}$ , and  $A$  and  $A^c$  are complementary. If all miners have the same chance of going offline, i.e.  $p_1 = p_2 = p_3 = \dots = p_{l_m} = p_d$ , the expectation of  $l_{on}$  is:

$$E[l_{on}] = l_m - E[X] = l_m(1 - p_d), \quad (3)$$

or equivalently,

$$l_m = \frac{E[X]}{p_d} = \frac{E[l_{on}]}{1 - p_d}. \quad (4)$$

We demonstrate how network churn affects the blockchain system performance in Section V-C.

2) *Mining Reward*: Bank determines the mining reward  $R$  according to the status of local economy and its budget. It becomes harder for each individual miner to find a valid block when there are more competitors (cf. assumption 1). As the expected reward reduces, mining becomes less profitable. We here calculate the maximum number of miners allowed to keep them incentivized. We denote the mining cost as  $\eta$  per hash, similar to the one in [28]. With individual hash rate being  $h$ , each miner's operational cost per second is  $\eta h$ . The expected hash cost of the entire network per second  $C_H$  when all miner are online (worst case) is  $C_H = l_m \eta h$ . Since miners are equal, their expected revenue per mining round is hence  $R' = R/l_m$ . We do not consider rewards provided for stale blocks (cf. assumption 3). We derive the expected profit per mined block for each individual as:

$$\Pi = R' - \frac{C_H T}{l_m} = \frac{R}{l_m} - \eta h T. \quad (5)$$

Clearly, for mining to be profitable,  $\Pi$  should be greater than 0, i.e.  $l_m$  and  $R$  should satisfy:

$$l_m < \frac{R}{\eta h T}. \quad (6)$$

3) *Minimal Connectivity Requirement*: A key factor of blockchains is the synchronization across the network, i.e. every node should obtain a copy of the ledger. Individual miners may want to reduce their data usage by connecting to fewer other nodes, but with security considerations all miners are encouraged to connect to as many other peers as possible. We here provide an intuitive picture on the minimal direct connections required for efficient information propagation across the network. For the simplicity of analysis, we assume on average every miner maintains connections with  $l_c$  other nodes,  $0 < l_c \leq l_m$ . Once Miner A receives a transaction or a block, she sends it to all the directly connected nodes. The same

procedure repeats until the message reaches every node in the network. We define information propagation from one network node to another as a hop. Restricting maximum number of hops to  $k$ , we obtain:

$$1 + l_c + l_c(l_c - 1) + \dots + l_c(l_c - 1)^{k-1} = 1 + l_c \sum_{i=0}^{k-1} (l_c - 1)^i \geq l_m. \quad (7)$$

Solving Equation 7 under given  $k$  and  $l_m$  (not  $l_{on}$  as eventually all miners will store the blockchain), we obtain  $\gamma$ , fraction of nodes that one should directly connect to as:

$$\gamma \geq \frac{\min(l_c)}{l_m}. \quad (8)$$

#### E. Summary

Although the bank has every intention to have fewer mining nodes to reduce the equipment cost, there should be sufficient miners to keep the fairness of the system and meet the security requirements. The bank should also combine real-world churn rate data and equipment costs to determine the number of mining nodes needed and device type according to particular use cases. Furthermore, all villagers, especially miners, are recommended to increase their connectivity for a better synchronization even though this leads to higher bandwidth usage.

### V. EVALUATION

In this section, we first validate our transaction processing model by comparing it to a real deployment of a private Ethereum blockchain and make predictions about the system behaviour based on the validated model. We also dimension the local blockchain network design by identifying the minimal average connection required for each mining node. Finally, we evaluate the blockchain performance under various *Network Delays* and *Network Churns* on the experimental deployment.

#### A. Experimental Environment

We have deployed a private Ethereum blockchain residing on multiple virtual machines running Ubuntu Linux v16.04 in an Openstack environment.<sup>6</sup> Each virtual machine is given 1 virtual CPU core, 2 GB of memory, and 10 GB of persistent storage to meet the minimum hardware requirement for running Ethereum. Elastic Search<sup>7</sup> was used to store block-related information. The network behaviour was monitored using the Python Web3 Library.<sup>8</sup>

All virtual machines are linked together in a low-latency local network that can be customized on demand. Our setup peers every node using a single Gigabit Ethernet switch and forms a star topology similar to the cellular network in Figure 1a. With this setup, a communication round-trip time between every two nodes is less than 1 ms on average. We employ Linux traffic control to introduce delays to emulate high latency in mobile networks and configure the Linux kernel firewall with Iptables<sup>9</sup> to emulate churns.

We have used Geth v1.6.4<sup>10</sup> for all of our empirical evaluations. Before the actual experimentation, we let the system stabilize for a few hours to obtain appropriate parameters of the genesis block in later runs. Based on our observation, we set the initial *nonce value*, *gas limit* and *difficulty* to 0x42, 0x08000000, 0x400000 respectively to make the system stabilize quickly. To form a fully connected P2P blockchain network, we disabled the auto-discovery feature supported by Geth and configured the overlay network manually. Finally, if not mentioned otherwise, all experiments involved 10 mining nodes and 10 light nodes.

#### B. Model Validation and Numerical Sensitivity Analysis

1) *Model Validation*: To validate our model in Section IV-B, we emulated a transaction processing scenario by issuing transactions at 1 tps (overall rate) from multiple clients to randomly selected addresses. We sort block timestamps in ascending order and calculate block generation time as the difference between two adjacent timestamps. To obtain transaction processing time, we record the timestamp upon initiation of a transaction and extract the inclusion timestamp from the transaction receipt. Figure 2a and Figure 2b present the simulation and emulation results of 500 consecutive blocks on block time and transaction processing time respectively. Both illustrate a strong agreement between modelling and experimental measurements. Additionally, an inter-comparison between Figure 2a and Figure 2b confirms that transaction processing and block generation are highly correlated.

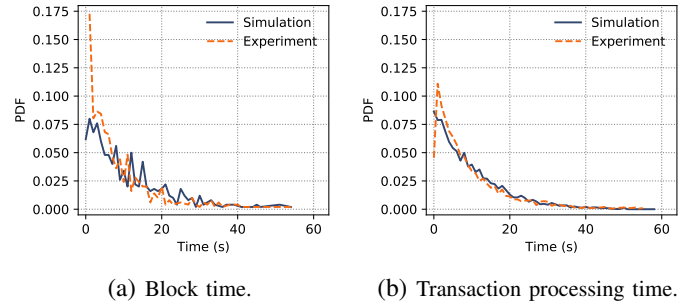


Figure 2: Model validation in terms of block generation and transaction processing at 1tps.

2) *Impact of Transaction Arrival Rate*: We then intended to empirically study the impact of transaction rate on transaction processing and benchmark the local blockchain performance. However, when increasing the arrival rate, we hit the transaction sending limit before the processing limit due to implementation bugs in the Ethereum version we used. Indeed, with more than 3 tps, the transaction generator stopped working due to connection errors as the algorithm thinks some nodes are flooding the system with too many requests. Despite the transaction generator worked smoothly under lower arrival rates, in the end only 4246 out of 17265 transactions were mined at 2 tps. We hence decided to analytically study the system behaviour under multiple transaction rates.

Using the validated model in Part V-B1, we further generated transactions at various rates including 0.2 tps, 1 tps, 5 tps, and 25 tps in our simulation. Figure 3a and Figure 3b compares the 50th, 70th, 90th, 95th and 99th percentiles of block time

<sup>6</sup><https://www.openstack.org>

<sup>7</sup><https://www.elastic.co>

<sup>8</sup><https://web3py.readthedocs.io/en/latest/index.html>

<sup>9</sup><https://help.ubuntu.com/community/IptablesHowTo>

<sup>10</sup><https://github.com/ethereum/go-ethereum>






Figure 3: Simulation with multiple transaction rates.

and transaction processing time under all the simulated arrival rates. These results confirm that with a sufficient block size, block creation time and transaction processing time are not affected by the arrival rate.

3) *Minimal Connection Requirement*: Based on the analysis presented in Part IV-D3, we obtain the minimal connection requirement when varying number of hops  $k$  from 1 to 4 in (7) with a total number of miners  $l_m$  increasing from 4 to 100. Figure 4 displays the results in the number of nodes and fraction of the network. As the left-hand side of (7) increases exponentially with  $l_c$ , minimum  $l_c$  that satisfies the equation may remain unchanged for a range of  $l_m$ , which results in steps or spikes on the curves. As shown in Figure 4a, when  $k = 2$ , with a total of 100 miners, each individual needs to maintain at least 10 connections. When  $k = 3$ , the minimum number of connections required reduces to 4. Generally speaking, if only one hop is allowed for information propagation, all miners need to connect with all other miners. While if more hops are allowed, miners can limit their connections to reasonably small values to save the bandwidth usage. This indicates the system’s ability to scale without miners exhausting their network resources.

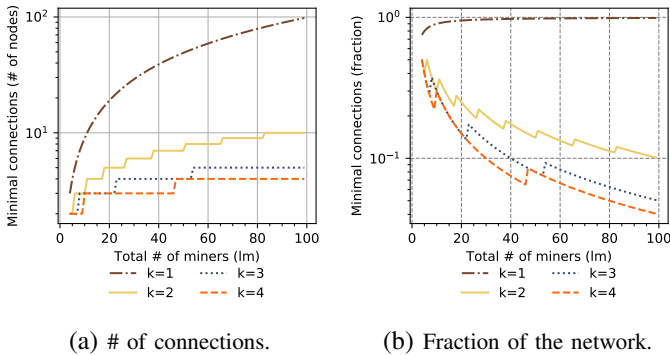


Figure 4: Minimal connection requirement.

### C. Effect of Network Disturbances

To further evaluate the performance of the local blockchain network, we investigate its behavior under disturbances such as *Network Delays* and *Network Churns*. Since it was impossible to obtain the whole picture from transaction processing (cf. Part V-B2), we have considered block generation time to be an alternative evaluation metric (cf. Part V-B1).

1) *The (non)Impact of Network Delays*: In order to understand the stability of our proposal when inter-node delay increases, we introduced various delays including 0, 10 ms, 50 ms, 100 ms, 500 ms, 1000 ms per connectivity channel.<sup>11</sup> Figure 5a shows the 50th, 70th, 90th, 95th, 99th percentiles of block time, from which we can see block time stays stable even with delays reaching 1 second. We ascribe this to PoW difficulty adjustment where network delay causes time difference between two adjacent blocks to increase, the internal algorithm reduces difficulty level to achieve a shorter block time. Figure 5b illustrates block difficulty under multiple network delays. A decreasing trend can be observed in the difficulty level when delay increases. This behaviour helps to maintain the transaction processing speed, but is “unhealthy” because with a fixed total network hash rate, a lower difficulty level leads to more stale blocks and inconsistencies.

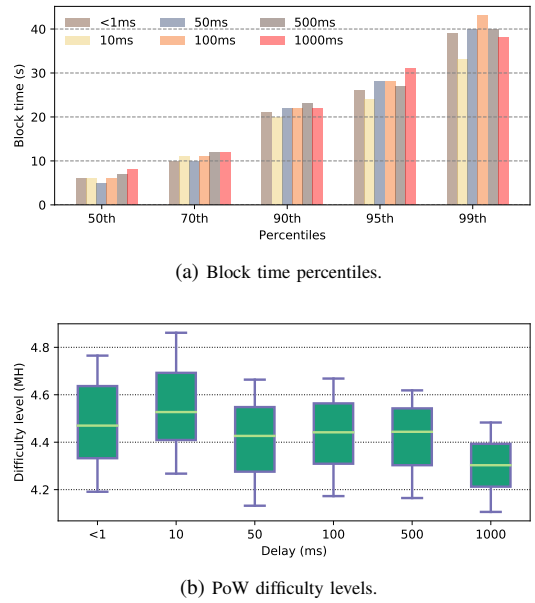


Figure 5: System behaviour with network delays.

### 2) Network Churns:

a) *Transient Response*: One of the main characteristic of P2P networks concerns the churn rate of their nodes. We emulated a scenario in which some miners go offline at time  $t_0$  and observed variation in block creation time. As shown in Figure 6, when some miners become offline, block time experiences a significant increase followed by a “resolving period” during which PoW difficulty is adjusted in response to changes in network hash rate. To find the end point of the “resolving period”, we select the timestamp from which the absolute variation is less than 5% for the next 3 consecutive points. Duration of “resolving period” is measured as 361 s, 781 s, 577 s, 527 s, 969 s respectively. Overall, it increases with the number of offline nodes. While the network recovers, we have also observed the behaviour of “backwards mining”, meaning that newer blocks sometimes have smaller block number than older ones. The disagreement is resolved with the stabilization of the network.

<sup>11</sup>Today’s mobile network delays can hardly go beyond 1 second: <https://hpb.co/mobile-networks/>

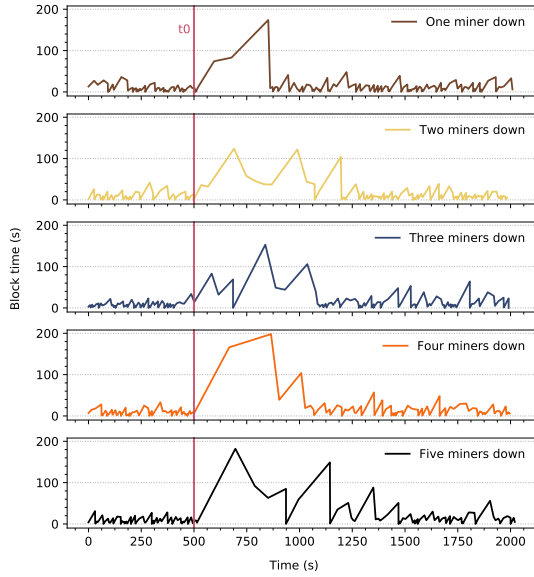


Figure 6: Transient response to mining nodes going offline.

*b) Changes Over Time:* We then tested system performance when each miner has a churn rate of 10%, 20%, 30%, 40% and 50%. We ran each set of experiment for approximately 2 hours and changed miner status (online or offline) every 20 minutes. Figure 7 shows the 50th, 70th, 90th, 95th, 99th percentiles of block generation time, from which we can see block time does not vary much across all churn rates. In other words, number of miners do not affect speed of block generation or transaction processing.




Figure 7: System steady-state behaviour with network churns.

*3) Summary:* Despite being unable to test the processing limitation due to Ethereum implementation flaws and obtain all transactional data experimentally, we managed to study the effect of network disturbances from block time. Difference between empty blocks and blocks with transactions mainly lies in block size, hence the additional transmission delay which is expected to have similar impacts to network delays. Results show that delays and churns may cause block generation time to change especially in the transient state, however, due to the difficulty adjustment they do not necessarily slow down block generation in the long run. Overall, the main observed drawback of the local blockchain is that disturbances may cause more temporary inconsistencies.

## VI. IMPLEMENTATION

In this section, we demonstrate the feasibility of the system design with a prototype implementation containing a private Ethereum blockchain and an intermittently connected bank full node. We tested the synchronization delay of the bank node when varying the bandwidth and disconnection duration. We also measured the usage of data, CPU, and battery of the mobile wallet app using DDMS<sup>12</sup> and Battery Historian.<sup>13</sup>

### A. Setup

Figure 8 illustrates the setup with a bank full node, 2 shops, and 3 regular users. Each shop runs a mining node and in addition, Shop 1 owns two full nodes and Shop 2 owns one full node. Regular User 3 operates a miner while the other two regular users are light mobile clients. We summarize the device types, capabilities and Geth versions in Table II.

Table II: Devices and their capabilities.

	Miner	Full node (including bank)	Light node
Device	Dell Latitude 6430u	Raspberry Pi 3	Google Pixel XL
# of nodes	3	4	2
OS	Ubuntu 17.04	Raspbian Jessie	Android v8.0.0
RAM	4GB	1GB	4GB
Geth	v1.6.5	v1.6.5	v1.6.7

We used a D-Link DSR-250N Wi-Fi router connecting to the Oulu public WAN network (PanOULU) [24] to represent the community base station. We used an IEEE 802.11n Wi-Fi network to interconnect the above nodes except the bank node which is directly connected to the PanOULU network and periodically joins the blockchain via the Wi-Fi router's public network interface. We configured the bank node's backhaul bandwidth via the router's WAN port. We used the auto-discovery protocol of Geth on all nodes.

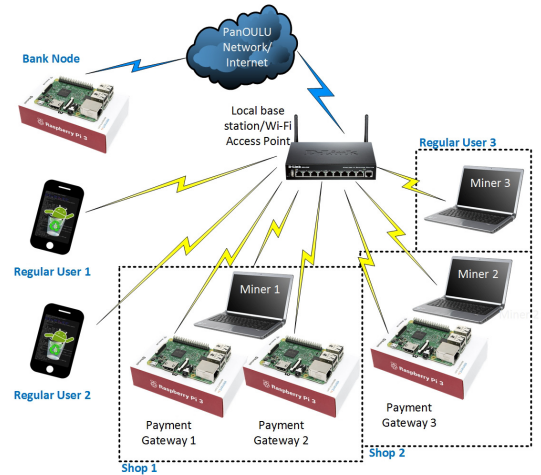


Figure 8: Prototype implementation.

We developed a smart contract as the wallet app in Solidity v0.4.12 to create fiat and digital user accounts as well as a special Token to perform functions including token supply, token conversion, token transferring and miner rewarding as described in Section III.

<sup>12</sup><https://developer.android.com/studio/profile/ddms.html>


<sup>13</sup><https://developer.android.com/studio/profile/battery-historian.html>



## B. Experiments

1) *Bank Node Synchronization Delay*: We measured bank synchronization delay (i.e. time taken to update the bank node) while increasing the backhaul bandwidth from 128 Kbps to 10 Mbps. We maintained a constant transaction rate of 2 tps and varied the disconnection time from 1 min to 10 min. Figure 9a illustrates the averaged synchronization delay and variation over 10 runs of each scenario, where synchronization delay is roughly proportional to the disconnection duration under each a fixed backhaul bandwidth.

2) *Mobile Wallet Performance*: We then measured data, CPU and battery used by the mobile wallet in its idle state (i.e. *Inactive* mode) and when it sends one transaction per minute (i.e. *Active* mode) for a total of one hour. Note that without sending any transactions, the mobile light client continuously synchronizes with the network and downloads block headers. As shown in Figure 9b, our wallet app uses only 2.71 MB in total in its idle state, and an additional 21.65 KB to perform one transaction per minute. The overall CPU usage increased from 51.9 s to 127.7 s within the one-hour period. When it comes to power consumption, the wallet used 3.2 mAh or 0.04% of the mobile phone's total battery usage in its idle state and 19.8 mAh or 0.11% when sending transactions. These results confirm our wallet app requires low bandwidth, CPU processing and power for its operations compared to available cryptocurrency clients [2].



(a) Bank node synchronization delay. (b) Mobile wallet data usage.

Figure 9: Test results on prototype implementation.

## VII. RELATED WORK

We categorize related work into i) Pervasive, Delay-Tolerant Networks, ii) Micro-Payment Systems, and iii) Blockchain Cost Models.

i) **Pervasive, Delay-Tolerant Networks**: Blattman et al. [3] found delay-tolerant networks are sufficient for digital services to meet the needs of most rural communities. Pentland et al. [25] later developed DakNet that uses a pervasive mobile coverage to enable asynchronous digital services in India and northern Cambodia. DakNet is remarkably low-cost and well-received by local users and is more accessible than a centralised, community telephone. The success of DakNet has proven that decentralization is an effective way to deliver service to remote regions.

ii) **Micro-Payment Systems**: A number of micro-payment schemes were proposed in the mid-to-late 1990s. Some of them leverage SMS or USSD of the cellular networks, for instance, the BAAC in Thailand [11] and the M-Pesa Service in Kenya [21]. However, SMS messages are easily spoofable and hence require additional user verifications for security,

and USSD could be affected by session time-outs. Other micro-banking systems are the early-generation token payment platforms secured by time-lock puzzles, e.g. PayWord and MicroMint [27]. IBM also proposed a Micro-iKP protocol for frequent micro-payments [13] using "coupons" sent between players and verified by the bank. This scheme however, increases the communication cost. Bitcoin [22], the first successful P2P cryptocurrency, was established in 2009 and has received great attention. Other alternative coins such as Ethereum [31] and Dogecoin were later developed. Although blockchain was originally designed to operate without centralised authorities, many major banks nowadays are also interested in using it to collaboratively settle transactions.<sup>14</sup>

iii) **Blockchain Cost Models**: Croman et al. [8] did a reality check of Bitcoin and analyzed the cost to confirm transactions. Rimba et al. [26] later proposed cost models using parameters such as *gas* and *gasPrice* to compare Ethereum with cloud services for business process execution. We stick to findings in [8] to obtain deployment and operation costs as the gas-related parameters can be controlled by the bank.

## VIII. DISCUSSION

By now we have demonstrated the design of a hybrid, blockchain-based payment scheme in an intermittently-connected environment. Although the ultimate control belongs to the bank, the reliable operation of the system is achieved via decentralization. Our system benefits all participants. Villagers are genuinely motivated to join as they can make low-fee, cash-less transactions efficiently. Miners are incentivised to work and receive rewards. By managing users' digital and fiat accounts, the bank can make a profit in a traditional way. While with a distributed, self-regulated blockchain for transaction processing, the bank can avoid single points of failure, and enjoys a lower deployment cost compared to building physical infrastructures and hiring experienced staff. However, some aspects of the system design can be further improved.

### A. Blockchain Inefficiency

PoW is a heavy consensus algorithm, and it becomes more inefficient in public cryptocurrencies as players compete against each other to earn profits. However, it does not appear to be a problem in our implementation with low-power mining nodes. One of the most promising replacements is the Proof-of-Stake (PoS) proposed by King et al. [18]. The idea is to assign mining rights based on the "coin-age", which equals the currency amount times its holding period. More recent proposals include Bitcoin-NG [10] that utilizes "leader selection" and "micro-blocks", LocalCoin [6] which operates on mobile devices without the Internet, and the Red Belly Blockchain [7] which incorporates a fast, leader-free Byzantine consensus. However, as PoS and other protocols also have vulnerabilities [15], we expect PoW to continue being the main consensus algorithm in blockchain-type systems in the near future.

### B. Security Considerations

1) *Temporary Inconsistency (not a real threat)*: Even though *blockchain forks* or *stale blocks* are an important indicator of inconsistency [9], if everyone behaves honestly, eventually all conflicts can be resolved. Hence, stale blocks themselves are

<sup>14</sup><https://www.reuters.com/article/us-blockchain-banks/six-big-banks-join-blockchain-digital-cash-settlement-project-idUSKCN1BB0UA>

not a direct threat to network security. However, they increase the chance of *double-spend* attacks that are usually caused by disagreements within the network [12].

**Countermeasure:** As suggested by Karame et al. [16], one way to mitigate *double-spend* attacks is to apply a “listening period” of a few seconds on the recipient side. This can be integrated into the mobile wallet design.

2) *Network Partitioning (a real threat):* Wust et al. [32] managed to perform an eclipse attack on a private ethereum blockchain and Natoli et al. [23] leveraged the network partitioning to perform the balanced attack.

**Countermeasure:** In our private blockchain setting, the permitted miners may turn malicious at a later time. To enhance security, we can enable the bank to deploy a smart contract that tracks node connections and detects network partitioning. We here present a simple method described as follows. We start off with a randomly selected miner  $m_1$  belonging to the set  $N = \{m_1, m_2, m_3, \dots, m_{l_m}\}$ . Miners that are connected directly to  $m_1$  belong to the set  $P_1 = \{m_2, m_3, m_4, \dots, m_k\}$ . For each element in  $P_1$ , we obtain similar sets of  $P_2, P_3$ , etc. This process is repeated until we find a set  $P_x \in P_1 \cup P_2 \cup \dots \cup P_{x-1}$ . If there are still remaining elements, i.e.  $P_x \subset N$ , a subgraph is detected. When the bank is connected, it accesses the results generated by this smart contract. The bank can then take action, e.g. to force nodes in the subgraph to restart and establish new connections. Note that this method does not prevent attacks or take effect in real-time.

## IX. CONCLUSION

In this paper, we propose a blockchain-based payment scheme in a remote region setting which has intermittent connectivity to the bank’s central system. To do so, we have introduced a novel way of using smart contracts by making them responsible for the admission control of the system, account management, mining rewards, and token creation. We developed mathematical models to analyze system dynamics and the costs for setup and operations. We validated the proposed models on a private Ethereum testbed and demonstrated the practicality of system design using laptops and mobile devices. Through a comprehensive study of the local blockchain operation, we showed its stability under network disturbances. We then demonstrated the whole system operates well in resource-constrained, dynamic, intermittently connected environments. Future work will address the security issues further, but we are confident that the security threats will not be a major bottleneck in realizing our proposal.

## REFERENCES

- [1] “KUHA Mobile Network,” <https://www.kuha.io>, 2017, [Online; accessed 19-September-2017].
- [2] “Requirements and warnings - bitcoin core,” <https://bitcoin.org/en/bitcoin-core/features/requirements>, 2017, [Online; accessed 21-September-2017].
- [3] C. Blattman, R. Jensen, and R. Roman, “Assessing the need and potential of community networking for developing countries: A case study from india,” *Harvard Center for International Development.*, 2002.
- [4] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, “Sok: Research perspectives and challenges for bitcoin and cryptocurrencies,” in *IEEE Symposium on Security and Privacy (SP)*, 2015.
- [5] M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan, “On the instability of bitcoin without the block reward,” in *Proc. of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [6] D. Chatzopoulos, S. Gujar, B. Faltings, and P. Hui, “Localcoin: An ad-hoc payment scheme for areas with high connectivity,” *arXiv preprint arXiv:1708.08086*, 2017.
- [7] T. Crain, V. Gramoli, M. Larrea, and M. Raynal, “(leader/randomization/signature)-free byzantine consensus for consortium blockchains,” *arXiv preprint arXiv:1702.03068*, 2017.
- [8] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. G. Sirer et al., “On scaling decentralized blockchains,” in *Financial Cryptography and Data Security*, 2016.
- [9] C. Decker and R. Wattenhofer, “Information propagation in the bitcoin network,” in *IEEE P2P 2013 Proceedings*. IEEE, 2013, pp. 1–10.
- [10] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, “Bitcoin-ng: A scalable blockchain protocol,” in *USENIX NSDI 16*.
- [11] D. Fitchett, “Bank for agriculture and agricultural cooperatives (baac), thailand (case study),” *Washington DC: Consultative Group to Assist the Poorest (CGAP) Working Group on Savings Mobilization*, 1999.
- [12] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, “On the security and performance of proof of work blockchains,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [13] R. Hauser, M. Steiner, and M. Waidner, *Micro-payments based on iKP*. IBM TJ Watson Research Center, 1996.
- [14] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, “Eclipse attacks on bitcoin’s peer-to-peer network,” in *USENIX Security Symposium*, 2015.
- [15] N. Houy, “It will cost you nothing to kill a proof-of-stake cryptocurrency,” *Browser Download This Paper*, 2014.
- [16] G. Karame, E. Androulaki, and S. Capkun, “Two bitcoins at the price of one? double-spending attacks on fast payments in bitcoin,” *IACR Cryptology ePrint Archive*, vol. 2012, no. 248, 2012.
- [17] S. Kasahara and J. Kawahara, “Priority mechanism of bitcoin and its effect on transaction-confirmation process,” *arXiv preprint arXiv:1604.00103*, 2016.
- [18] S. King and S. Nadal, “Ppcoin: Peer-to-peer crypto-currency with proof-of-stake,” *self-published paper*, August, vol. 19, 2012.
- [19] T. J. MacDonald, D. W. Allen, and J. Potts, “Blockchains and the boundaries of self-organized economies: predictions for the future of banking,” in *Banking Beyond Banks and Money*. Springer, 2016.
- [20] A. Mackenzie, “The fintech revolution,” *London Business School Review*, vol. 26, no. 3, pp. 50–53, 2015.
- [21] I. Mas and D. Radcliffe, “Mobile payments go viral: M-pesa in kenya,” 2010.
- [22] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008.
- [23] C. Natoli and V. Gramoli, “The balance attack against proof-of-work blockchains: The r3 testbed as an example,” *arXiv preprint arXiv:1612.09426*, 2016.
- [24] T. Ojala, J. Orajärvi, K. Puhakka, I. Heikkinen, and J. Heikka, “panoulu: Triple helix driven municipal wireless network providing open and free internet access,” in *Proceedings of the 5th International Conference on Communities and Technologies*. ACM, 2011, pp. 118–127.
- [25] A. S. Pentland, R. Fletcher, and A. Hasson, “Daknet: Rethinking connectivity in developing nations,” *Computer*, vol. 37, no. 1, pp. 78–83, Jan. 2004.
- [26] P. Rimba, A. B. Tran, I. Weber, M. Staples, A. Ponomarev, and X. Xu, “Comparing blockchain and cloud services for business process execution,” in *ICSA*, 2017.
- [27] R. L. Rivest and A. Shamir, “Payword and micromint: Two simple micropayment schemes,” in *International Workshop on Security Protocols*. Springer, 1996, pp. 69–87.
- [28] P. R. Rizun, “A transaction fee market exists without a block size limit,” 2015.
- [29] R. C. Shah, S. Roy, S. Jain, and W. Brunette, “Data mules: modeling a three-tier architecture for sparse sensor networks,” in *Proc. of IEEE International Workshop on Sensor Network Protocols and Applications.*, May 2003.

- [30] Y. Sompolinsky and A. Zohar, "Accelerating bitcoin's transaction processing. fast money grows on trees, not chains." *IACR Cryptology ePrint Archive*, vol. 2013, no. 881, 2013.
- [31] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, 2014.
- [32] K. Wüst and A. Gervais, "Ethereum eclipse attacks," ETH Zurich, Tech. Rep., 2016.