An Ultra-Scalable Blockchain Platform for Universal Asset Tokenization: Design and Implementation

Ahto Buldas[1,4] (✉), Dirk Draheim[2] (✉), Mike Gault[3], Risto Laanoja[4,1], Takehiko Nagumo[5,6], Märt Saarepera[7,4], Syed Attique Shah[8], Joosep Simm[4], Jamie Steiner[4], Tanel Tammet[9], and Ahto Truu[4]

[1] Centre for Digital Forensics and Cyber Security, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia
ahto.buldas@taltech.ee

[2] Information Systems Group, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia
dirk.draheim@taltech.ee

[3] Guardtime, Avenue d'Ouchy 4, 1006 Lausanne, Switzerland
mike.gault@guardtime.com

[4] Guardtime, A. H. Tammsaare tee 60, 11316 Tallinn, Estonia
{risto.laanoja,mart.saarepera,joosep.simm,jamie.steiner,ahto.truu}@guardtime.com

[5] Mitsubishi UFJ Research and Consulting, 5-11-2 Toranomon, Minato City, Tokyo, 105-8501, Japan
takehiko.nagumo@murc.jp

[6] Graduate School of Management, Kyoto University, Yoshida Honmachi, Sakyo-ku, Kyoto-shi Kyoto 606-8501, Japan

[7] Martest Research, Pärnu mnt 160a, 11317 Tallinn, Estonia

[8] School of Computing and Digital Technology, Birmingham City University, Millennium Point, B4 7XG, Birmingham, United Kingdom
Syed.Shah2@bcu.ac.uk

[9] Applied Artificial Intelligence Group, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia
tanel.tammet@taltech.ee

Abstract. Since its introduction with Bitcoin in 2009, blockchain technology has received tremendous attention by academia, industry, politics and media alike, in particular, through extended blockchain-based visions such as smart contracts, decentralized finance, and, most recently, Web3. The critical prerequisite for any such blockchain-based vision to be turned into reality is uncapped scalability. Furthermore, and equally important, blockchain technology needs to transcend the stage of specialized tokens into an adaptive, heterogeneous tokenization platform. In this paper, we explain the Alphabill family of technologies that addresses both unlimited scalability and unrestricted adaptivity. We deliver a sharded blockchain technology with unlimited scalability and performance, called KSI Cash, which is based on a new form of electronic money scheme, the bill scheme. We present performance tests of KSI Cash that we have conducted with the European Central Bank and a group of eight national central banks from the Eurosystem in order to assess the technological feasibility of a digital euro, showing the system operating with 100 million wallets and 15 thousand transactions per second (under simulation of realistic usage), having an estimated carbon footprint of 0.0001g CO2 per transaction (Bitcoin = 100 kg and more); furthermore, showing the system operating with up to 2 million payment orders per second, an equivalent of more than 300,000 transactions per second (in a laboratory setting with the central components of KSI Cash), scaling linearly in terms of the number of deployed shards. We explain, in detail, the key concepts that unlock this performance (i.e., the concepts of the bill money scheme). The results provide evidence that the scalability of our technology is unlimited in both permissioned and permissionless scenarios, resulting into the Alphabill Money technology. Next, we contribute the architecture of a universal tokenization platform that allows for universal asset tokenization, transfer and exchange as a global medium of exchange, called Alphabill platform. We reveal the crucial conceptual and technical contributions of the platform's architecture and their interplay, including the data structures of KSI Cash and Alphabill Money, the dust collection solution of Alphabill Money, and the atomic swap solution of the Alphabill platform.

Keywords: Blockchain · Bitcoin · cryptocurrency · sharding · digital euro · decentralized finance · DeFi · Web3 · asset tokenization · KSI Cash · Alphabill

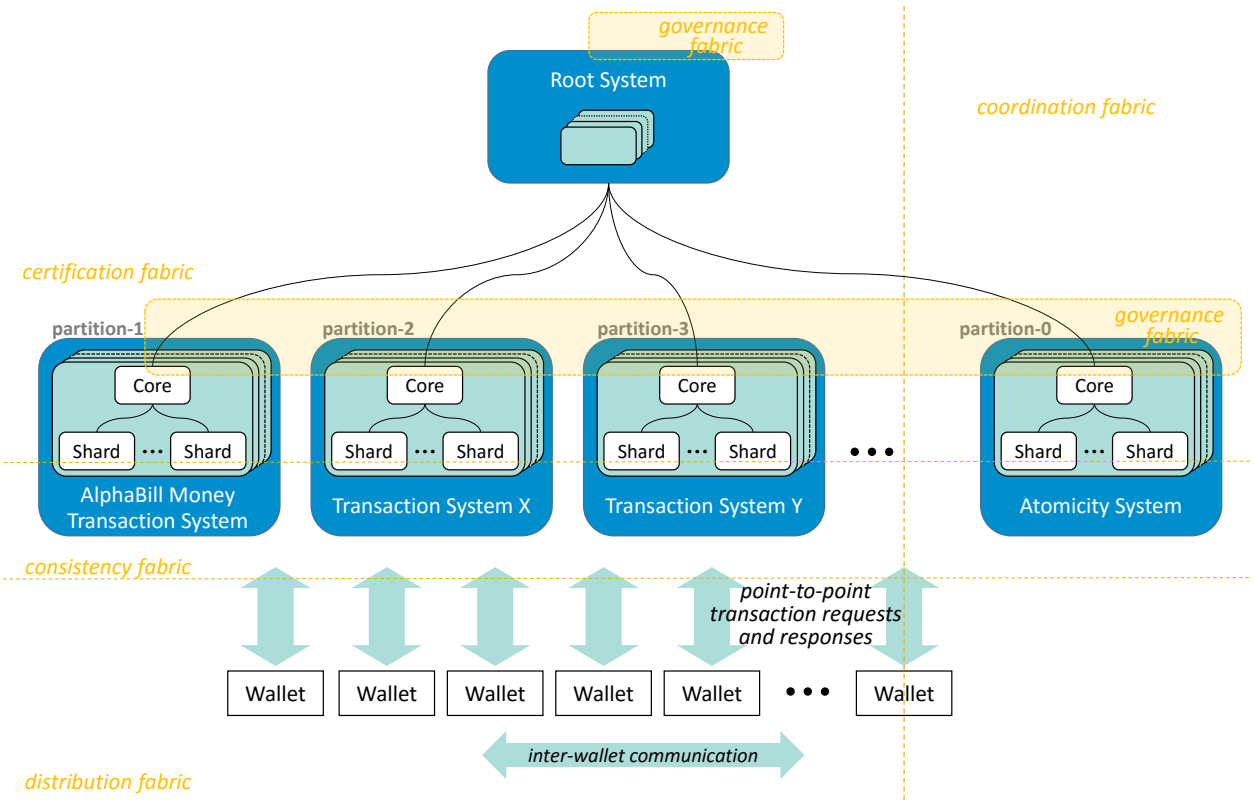

Fig. 1. Alphabill high-level architecture.

1 Introduction

Blockchain technology has received tremendous attention from both industry and academia in the last decade. Since the introduction of blockchain technology with the cryptocurrency Bitcoin [1] in 2009, we have seen a plethora of cryptocurrencies and blockchain-based solutions. In connection with that, we have seen the emergence of a series of extended blockchain-based visions such as smart contracts [2,3,4,5], decentralized finance (DeFi) [6], and, most recently, Web3.

The prerequisite for any valuable blockchain-based vision to be turned into reality is uncapped scalability. With an average transaction rate of 7 transactions per second, Bitcoin is far from the transaction rates of today's established payment systems such as VisaNet, which has a peak performance of 24,000 transactions per second according to [7]. According to [8], there has been a total of approx. 435 billion annual USD transactions in 2017 via the three biggest players in the field alone, i.e., Visa (300 billion), MasterCard (75 billion), and UnionPay (60 billion), which amounts to an *averaged* transaction rate of approx. 13,800 transactions per second. Therefore, also blockchain technologies need to be capable of thousands of transactions per second. If micro-payments and related visions such as the Internet of Things (IoT) are to be enabled by blockchain technology [9,10,11,12,13], much higher transaction rates will be needed already in the near future, i.e., in the range of hundreds of thousands of transactions per second and even way beyond. Also, it is often overlooked that shorter and shorter transaction durations (transaction settlement times) will become increasingly important against the background of these developments.

As a next problem, tokenization platforms seen so far often suffer a reality check, by either oversimplifying today's established monetary system [14] or neglecting the reality of today's institutional stack [14,15] and (inter-)/organizational settings [16] or both; and, therefore, lack completeness of vision. Typically, new

tokenization solutions proclaim disruptiveness in their field (at different levels: organizations, business domains, the whole monetary system, or the whole society), however, each of the single tokenization solutions is, itself, typically not prepared for *adaptation*, i.e., vulnerable to changes in its environment. Therefore, having more and more specialized tokenization solutions will no longer advance the field any more in the future. What is needed instead, is an *adaptive*, heterogeneous tokenization platform that allows for launching and coordinating individual blockchain technologies in a systematic manner.

In this paper, we present the Alphabill family of technologies that addresses both *unlimited scalability* as well as *unrestricted adaptivity*. We make the following contributions:

- *KSI Cash.* We deliver a blockchain technology with essentially unlimited scalability and performance, called KSI Cash. KSI Cash has been developed for experiments of the European Central Bank and a group of eight central banks from the Eurosystem[10] (Estonia, Germany, Greece, Ireland, Italy, Latvia, Spain, the Netherlands) to assess the feasibility of a digital euro [17]. KSI Cash has been developed as a joint design science effort of the Estonian Central Bank (Eesti Pank) and Guardtime [18,19]. The scalability of the Alphabill platform is founded on the scalability of its individual partitions. KSI Cash is part of this foundation, and a part of the Alphabill family of technologies.
- *Exhaustive KSI Cash Performance Evaluation.* We present diverse performance tests of KSI Cash that we have conducted with the European Central Bank and the said group of eight central banks [17]. We show the system operating with 100 million wallets and 15 thousand transactions per second under simulation of realistic usage. Furthermore, we show the system operating with up to 2 million payment orders per second (an equivalent of more than 300,000 transactions per second in our technology) in a laboratory setting that runs with the central components of KSI Cash. These tests show that the technology scales linearly in terms of the number of deployed shards. Furthermore, we are able to estimate the carbon footprint of KSI Cash as 0.0001g CO2 per transaction, again under the realistic usage scenario (as compared to Bitcoin, which amounts to much more than 100 kg CO2 per payment transaction according to several existing studies).
- *Bill Scheme.* We explain, in detail, the key concepts that unlock the performance of KSI Cash, i.e., a new form of electronic money scheme, the *bill scheme* [20,19], which achieves full decomposability by circumvention of monetary splits and joins. The analysis indicates that the performance of our technology is unlimited for practical purposes, in both permissioned and permissionless scenarios.
- *Alphabill Platform Architecture.* We provide the architecture of the Alphabill platform, see Fig. 1. The Alphabill platform enables universal asset tokenization, transfer and exchange as a global medium of exchange. The platform allows users to launch an arbitrary number of new so-called *partitions*. Each partition realizes an individual token and corresponding transaction system. The Alphabill platform provides the necessary innovations and protocols, along with respective languages, libraries and toolkits, to implement and launch partitions in such a way, that individual partitions as well as their coordination and interaction show unlimited scalability. We specify: the platform's elements, asset presentations, ledger certificates, transaction orders, sharding schemes and transaction system specifications.
- *Alphabill Money.* We introduce the Alphabill Money technology, which forms the native currency of the Alphabill platform. With Alphabill Money, we contribute an *extended bill scheme* that balances monetary splits and limited forms of monetary joins with a *dust collection* solution so that scalability is preserved. We specify the data structures of Alphabill Money.
- *Alphabill Money Dust Collection.* We specify the *dust collection* solution of Alphabill Money. We introduce and elaborate the notion of bill swap scenario and specify the dust collection process.
- *Alphabill Platform Atomicity Partition.* A key contribution is the design of a scalable *multi-asset atomic swap* solution (the atomicity system in Fig. 1) that enables cross-partition transactions that are needed, e.g., for financial transactions (parallel synchronized transfers of monetary and non-monetary assets). We specify the solution by defining the data structures of the atomicity system and defining a novel 3-phase-commit protocol.

[10] `https://www.ecb.europa.eu/ecb/orga/escb/`

With respect to usage of blockchain terminology we refer to the ISO standard ISO 22739 [21] and the NIST report NISTIR 8202 [22].

We proceed as follows. In Sect. 2, we discuss related work. In Sect. 3, we explain the architecture of the Alphabill platform and its design principles. In Sect. 4, we describe the technical principles and data structures of the central bank digital currency KSI Cash. In Sect. 5, we explain the design of the KSI Cash test bench and present the outcomes of a series of diverse test runs. In Sect. 6, we delve into the details of a series of essential Alphabill platform components, i.e., fundamental design elements, the Alphabill Money partition and the Alphabill atomicity partition. We finish the paper with a conclusion in Sect. 7.

2 Related Work

Scalability is a central concern for blockchain technology. In [23], the authors provide a survey of relevant approaches to improve the scalability of the Bitcoin network, including tuning Bitcoin protocol parameters, basic off-chain payment channels, duplex micro-payment channels [24], Bitcoin lightning channels [25], and off-chain payment networks. In [26], the authors provide an overview of scalability of blockchain systems and technologies for scalable blockchain systems. They discuss tuning of blockchain parameters, off-chain transactions, decoupling blockchain management from execution, and sharding. They identify Elastico [27] and OmniLedger [28,29] as sharding blockchain technologies.

2.1 Blockchain Sharding Approaches

Two most recent surveys of sharding in blockchains are provided by [30] and [31]. The survey in [30] provides an analysis of sharded blockchain technologies including aspects such intra-consensus settings, design of cross-shard atomicity, corresponding overhead, latency and theoretical throughput. The paper reviews Elastico [27], Chainspace [32,33], OmniLedger [28,29], RapidChain [34], and Monoxide [35] and investigates Ethereum 2.0 [36]. The survey in [31] is a survey of scaling blockchains; however, it focuses on sharding: "In particular, we focus on sharding as a promising first layer solution to the scalability issue" [31]. The survey in [31] analyses technologies in terms of committee formation and intra-committee consensus. In addition to the work analyzed by [30], the survey in [31] reviews [37], SSChain [38] and Ostraka [39] and, furthermore, considers the technology proposals Zilliqa[11], Harmony[12], Logos [13], and Stegos[14].

Elastico [27] is an early sharded, proof-of-work-based blockchain technology that did not yet provide cross-shard atomicity control.

OmniLedger [28,29] is a permissionless, sharded, UTXO-based blockchain technology that relies on a cross-shard, client-driven two-phase commit protocol (called Atoms) for the needed atomicity control. OmniLedger builds on the consensus protocol ByzCoin [40] by introducing sharding to it, resulting into the OmniLedger's consensus protocol ByzCoinX. Furthermore, OmniLedger builds on Hybrid Consensus [41,42] and Elastico [27]. Hybrid consensus [41,42] relies on *proof of work* and consumes a permissioned BFT (Byzantine fault tolerance) [43] protocol as a building block. To support power efficiency by proof-of-stake rather than proof-of-work consensus, OmniLedger builds on Ouroboros [44] and Algorand [45]. OmniLedger uses the RandHound protocol [46] targeting a scalable and at the same time bias-resistant sampling of representative validators.

Chainspace [32,33] is a sharded blockchain technology that supports user-defined smart contracts. The system exploits an implementation of PBFT (Practical Byzantine Fault Tolerance) [47], i.e., MOD-SMART (Modular State Machine Replication) [48], for intra-shard consensus. Cross-shard consistency is guaranteed by a distributed commit protocol, called S-BAC (Sharded Byzantine Atomic Commit). As opposed to OmniLedger's client-driven commit protocol, in S-BAC, the entire shard acts as a coordinator – rather than a single untrusted client.

[11] `https://docs.zilliqa.com/whitepaper.pdf`
[12] `https://harmony.one/whitepaper.pdf`
[13] `https://logos.network/whitepaper.pdf`
[14] `https://stegos.com/docs/stegos-whitepaper.pdf`

Ethereum upgrades[15] represents the current roadmap of Ethereum [49] as the successor of Ethereum 2.0 [36]. The roadmap aims at bringing proof-of-stake consensus and sharding to Ethereum via a component called *beacon chain*[16], which has already been described as part of Ethereum 2.0 [36]. "The Eth2 execution sharding (formerly known as Eth2 Phase 2) consists of one *beacon chain* and multiple *shard chains*. [...] The beacon chain mediates cross-shard communications." [50]

RapidChain [34] is a permissionless, sharded, UTXO-based blockchain. In RapidChain, shards are called *committees*. RapidChain builds on [51,52] to reduce communication overhead and latency in each committee. RapidChain aims at efficient cross-shard transaction verifications by employing a routing mechanism inspired by Kademlia [53], a peer-to-peer distributed hash table, for committees to discover each other.

Monoxide [35] is a sharded blockchain technology that is based on an *account money scheme*, i.e., an account/balance model, compare with [20]. Shards are called *zones* in Monoxide. For cross-shard transactions, Monoxide introduces a notion of *eventual atomicity* and suggests a lock-free cross-shard transaction design to achieve eventual atomicity, which aims at saving the overhead of a usual two-phase commit protocol needed to achieve *immediate atomicity*.

In [37], a permissioned, sharded, UTXO-based blockchain technology is suggested that aims at performance optimizations of the consensus protocol running within each individual shard, an efficient shard formation protocol, and a secure distributed (Byzantine fault tolerant) transaction protocol for cross-shard, distributed transactions (cross-shard commit protocol). The paper [37] considers technology in the context of trusted execution environments (TEEs).

SSChain [38] is a sharding-based protocol that addresses the problem of extensive data migration due to reshuffling of the network. For that, SSChain suggests a non-reshuffling structure in service of transaction sharding and state sharding. In order to achieve the elimination of data migration overhead, SSChain allows nodes to "freely join in one or more shards without reshuffling" [38]. The implementation of SSChain is based on the Bitcoin source code, therefore, SSChain uses the UTXO money scheme.

In [54], it is suggested to analyze the security of sharded blockchains (such as OmniLedger [28,29] and RapidChain [34]) by estimating the failure probabilities of single sharding rounds on the basis of failure probabilities of all shards. For this purpose, a hypergeometric distribution model is assumed. The proposed methodology is tried out and numerically analyzed on the basis of a large-scale trial.

Ostraka [39] is different from sharding-based blockchain protocols that split the network into shards, i.e., Ostraka shards the nodes themselves. The rationale behind the design of Ostraka is to prevent such denial-of-service attacks where a single shard is overloaded with the goal to bring down the complete transaction system.

In [55], the Mixed Byzantine Fault Tolerance (MBFT) protocol is suggested that uses so-called layering to separate specific node functions, i.e., separating the verification function from the demodulation process. Layering aims at reducing the load of nodes and this way improving the efficiency of the overall consensus mechanism. Layering is combined with sharding to further improve the protocol efficiency.

ZyConChain [56] is a sharded blockchain technology that follows the approach of a *general data model*, i.e., aims at being not restricted to a certain scheme such as UTXO. ZyConChain suggests to introduce three kinds of blocks in each shard, i.e., parent blocks, side blocks and state blocks, where each kind of block is maintained via its own consensus mechanism. Parent blocks are maintained via proof-of-work consensus, whereas for side blocks, a new consensus mechanism based on speculative Byzantine fault tolerance [57] is suggested.

In [58], an optimized data storage model for sharding-based blockchain technology is suggested. The approach suggests to exploit the Learning Machine (ELM) algorithm to identify *hot blocks* to be stored and queried locally according to several objective criteria, including historical popularity, hidden popularity and storage requirements of the block.

In [59], an adaptive resource-allocation algorithm is designed for optimizing transaction assignments in sharded, permissioned blockchains. The algorithm is designed on the basis of the drift-plus-penalty (DPP) technique [60] in the framework of stochastic network optimization [60,61].

[15] `https://ethereum.org/en/upgrades/`

[16] `https://github.com/ethereum/consensus-specs/blob/dev/specs/phase0/beacon-chain.md`

2.2 Central Bank Digital Currencies

RSCoin [62] is a central bank digital currency. Its architecture is based on shards (called *mintettes*) and a trusted central component (called *central bank*). Each mintette serves a subset of public client addresses, with a certain level of replication needed for majority voting. The mintettes are authorized by the central component and work together in creating consensus on a next block to be certified by the central component. The necessary communication is realized indirectly by the wallets via a two-phase-commit protocol, where each wallet makes its decision individually based on incoming majority votes.

The Hamilton Project [63], conducted by the Federal Reserve Bank of Boston together with the Massachusetts Institute of Technology Digital Currency Initiative, is a concept study on implementing central bank digital currency. In the project, they investigate two different architectures. The first one, called the *atomizer architecture* is a blockchain solution that relies on sharded transaction verification. The architecture follows the UTXO scheme. The central component of the architecture is the atomizer, which is responsible for collecting verified payments from shards and creating blocks, introducing an essential bottleneck to the overall system. The study proceeds with comparing the atomizer architecture with the so-called *2PC architecture*, which represents an instance of established (non-blockchain) transaction system technologies as found in today's banking domain.

2.3 A Large-Scale Off-Chain Payment Approach

An example of a large-scale off-chain payment solution is SLIP (Secure Large-Scale Instant Payment) [64]. SLIP uses an aggregate signature scheme to connect off-chain channels that enables locked-in tokens to circulate. This way, SLIP aims at improving the negotiability of locked-in tokens in off-chain channels. The generic construction of SLIP aims at defending against double-spending, over-spending, and malicious settlement attacks.

2.4 A Heterogeneous Multi-Chain Approach

Polkadot[17] is designed as a *heterogeneous multi-chain*. Polkadot distinguishes between the *relay chain* and *parachains* ("parallelised" chains) [65]. Polkadot itself provides the relay chain, which may host "validatable, globally-coherent dynamic data-structures" [65] called parachains, which do not necessarily have to be blockchains. With the parachains, Polkadot aims at blockchain scalability: "In principle, a problem to be deployed on Polkadot may be substantially parallelized – scaled out – over a large number of parachains" [65]. At the same time, parachains aim at blockchain heterogeneity: "In other words, Polkadot may be considered equivalent to a set of independent chains (e.g. the set containing Ethereum, Ethereum Classic, Namecoin and Bitcoin) [...]" [65]. In [65], Polkadot suggests a fully asynchronous interchain communication, i.e., parachains can dispatch transactions in other parachains or the relay chain. The dispatched transactions are fully asynchronous, i.e., "[...] there is no intrinsic ability for them to return any kind of information back to its origin" [65]. The communication mechanism is based on queuing: "Interchain transactions are resolved using a simple queuing mechanism based around a Merkle tree to ensure fidelity. It is the task of the relay-chain maintainers to move transactions on the output queue of one parachain into the input queue of the destination parachain." [65]

3 The Alphabill Platform

Alphabill is a system for trading with various kinds of digitalized (tokenized) assets. Alphabill enables swap transactions between different kinds of assets without trusted intermediaries, and we call those transactions *multi-asset swap transactions* (similar to multi-asset strategies, which also consist of various kinds of assets[18]).

[17] `https://polkadot.network/`
[18] `https://www.blackrock.com/us/individual/education/multi-asset-strategies`

Definition 1 (Multi-Asset Swap Transaction). *A* multi-asset swap transaction *is a swap transaction between different kinds of (tokenized) assets.*

The Alphabill system has its native electronic currency called Alphabill Money. With multi-asset swap transactions, digitized assets can be bought and sold for the native currency or swapped directly without involving the native currency. Multi-asset swap transactions are executed automatically – defined by blockchain rules, so that no nobody can break their trading contract agreements.

Each kind of asset is realized by a corresponding transaction system that is deployed to the Alphabill platform as so-called *partition*. Therefore, multi-asset swap transaction are also called cross-partition transactions in this paper.

3.1 The Alphabill Platform Architecture

The Alphabill platform, see Fig. 1, allows for the deployment of arbitrary many partitions, where each partition is a blockchain implementation of a transaction system. There are three specialized partitions that genuinely belong to the platform:

- *Root partition.* The root partition (or *root system*) keeps track of the registered transaction systems. The entities of the root partition are the registered transaction systems, each described by the system identifier α and the state tree type (see Def. 7). The root partition periodically receives the roots of the state trees of all transaction system and creates *uniqueness certificates* for the ledgers of the transaction systems. By uniqueness certificate, we mean any proof of uniqueness of the ledger such as proof of work, proof of stake, or proof of authority. The definite specifications of the root partition and the uniqueness certificate depend on the consensus protocol of the root partition, the particular choice of which is out of the scope of this paper.
- *Atomicity partition.* The atomicity partition (or *atomicity system*) supports atomic multi-asset swap transactions. The entities of the atomicity partition are multi-asset swap transactions, where each transaction is described by the identifier ι, the terms, and the status (complete/incomplete) of the transaction. The atomicity partition is described in detail in Sect. 6.3.
- *Alphabill Money partition.* The Alphabill Money partition (or *Alphabill Money system*) provides the default *medium of exchange* of the platform. The implementation of the money partition is described in Sect. 6.2

The communication protocols in the Alphabill platform (Fig. 1) can be divided into five logical groups (*fabrics*):

- The *consistency fabric*, which establishes the rules of communication between client applications (*wallets*) and the transaction systems, as well as the reactions of each transaction system to the communication, i.e., how the transactions (enacted by wallets) change the state of the transaction system.
- The *certification fabric*, which establishes the rules of communication between the transaction systems and the root system, as well as how the root system creates uniqueness certificates for all the transaction systems.
- The *distribution fabric*, which describes the rules of wallet-to-wallet communication.
- The *coordination fabric*, which describes the protocols between wallets and the atomicity system as well as the reaction of the atomicity system to protocol messages.
- The *governance fabric*, which describes the communication with the root system for registering, changing, and removing transaction systems from the framework.

3.2 The Alphabill Design Principles

The design goals of Alphabill are maximum security, maximum scalability and system viability [66]. Accordingly, the design principles of Alphabill are:

– *Secure-by-Design.* Alphabill is a partitioned blockchain system that operates according to abstract rules that cannot be broken unless the blockchain technology itself fails. Every transaction system in Alphabill has its specific rules. All communication between partitions is certified. Cross-partition messages have ledger certificates. A ledger certificate is a proof that the content of an inter-partition message is included in the ledger of the sender partition. The types of ledger certificates available in the Alphabill platform are described in Sect. 6.1.
– *Scalable-by-Design.* The Alphabill framework provides tools for designing transaction systems in a way that the ledger of a system is dynamically decomposable into autonomous *shards*. Each asset entity has its own independently verifiable sub-ledger.
– *Robust-by-Design.* The Alphabill framework provides tools for designing highly redundant transaction systems, following proven patterns of blockchain networks with many nodes.
– *Viable-by-Design.* The Alphabill framework is a self-adaptive system. It changes depending on the requirements of its business environment and the society [67,68,66,69,70,71,72,73,74]. This is only possible because Alphabill is a partitioned blockchain solution, where new transaction systems can be included, existing transaction systems can be modified, and transaction systems that are no longer needed can be removed from the system. This is not possible in monolithic permissionless blockchains such as Bitcoin, because in these blockchains all rules are fixed from the launch of the blockchain and henceforth cannot be changed.

4 KSI Cash Central Bank Digital Currency

KSI Cash [18,19,75] is a central bank digital currency (CBDC). KSI Cash represents an option of how to implement a currency partition of the platform in case it has to be controlled by a central bank.

KSI Cash is realized on the basis of a bill money scheme [20,19], or just bill scheme for short. The key property of any bill scheme is that it allows for sharding without atomic cross-shard transactions. The bill scheme mimics traditional cash. Traditional cash scales, without any automation, and, in a certain sense, the bill scheme can be considered as a direct digitization of traditional cash. In this paper we present a proof-of-concept of KSI Cash that has been implemented to evaluate its performance. This concept study is based on a pure bill money scheme as described in Sect. 4.1.

4.1 KSI Cash Data Structures

Bills and Payment Orders Figure 2 shows the data structures of KSI Cash consisting of the public blockchain data and the maintained system state. The digital asset of a bill scheme is called a *bill*. A bill has an owner and a value. During payments, it is possible to change the owners of bills, but it is not possible to change their values. This means that the value of a bill cannot be changed and the number and identities of all bills are fixed from the genesis of the ledger during all of its lifetime, i.e., bills cannot be destroyed, nor dynamically created. From that, it follows immediately, that there are no splits and joins of digital assets, as we have them, e.g., in Bitcoin UTXOs (unspent transaction outputs) [1,76]. Later, in implementations of KSI Cash in the field, the rigid change policy of bills needs to be loosened to allow for limited forms of splits and joins. The design solution is exactly about allowing for in-shard splits and joins, but still avoiding atomic cross-shard transactions.

Figure 2 shows a small example system with sixteen bills B_0 through B_{15}. Given a bill B_i, we say that i is the *identifier* (or *number*) of the bill. The bills are distributed over four shards $shard_1$ to $shard_4$ so that bills B_0 through B_3 are maintained by $shard_1$, bills B_4 through B_7 are maintained by $shard_2$ and so forth. In Fig. 2, time flows from left to right, divided into epochs (also called rounds), i.e., rounds of block creation, with n denoting the most recent epoch number. In a given epoch k, each payment order, denoted by B_i^k, that is successfully integrated into the block, changes the owner of the bill B_i (and, in the extended bill scheme, potentially also the value of that bill). A payment order B_i^k is a record

$$B_i^k = \langle i, owner(B_i^k), value(B_i^k) \rangle, \tag{1}$$

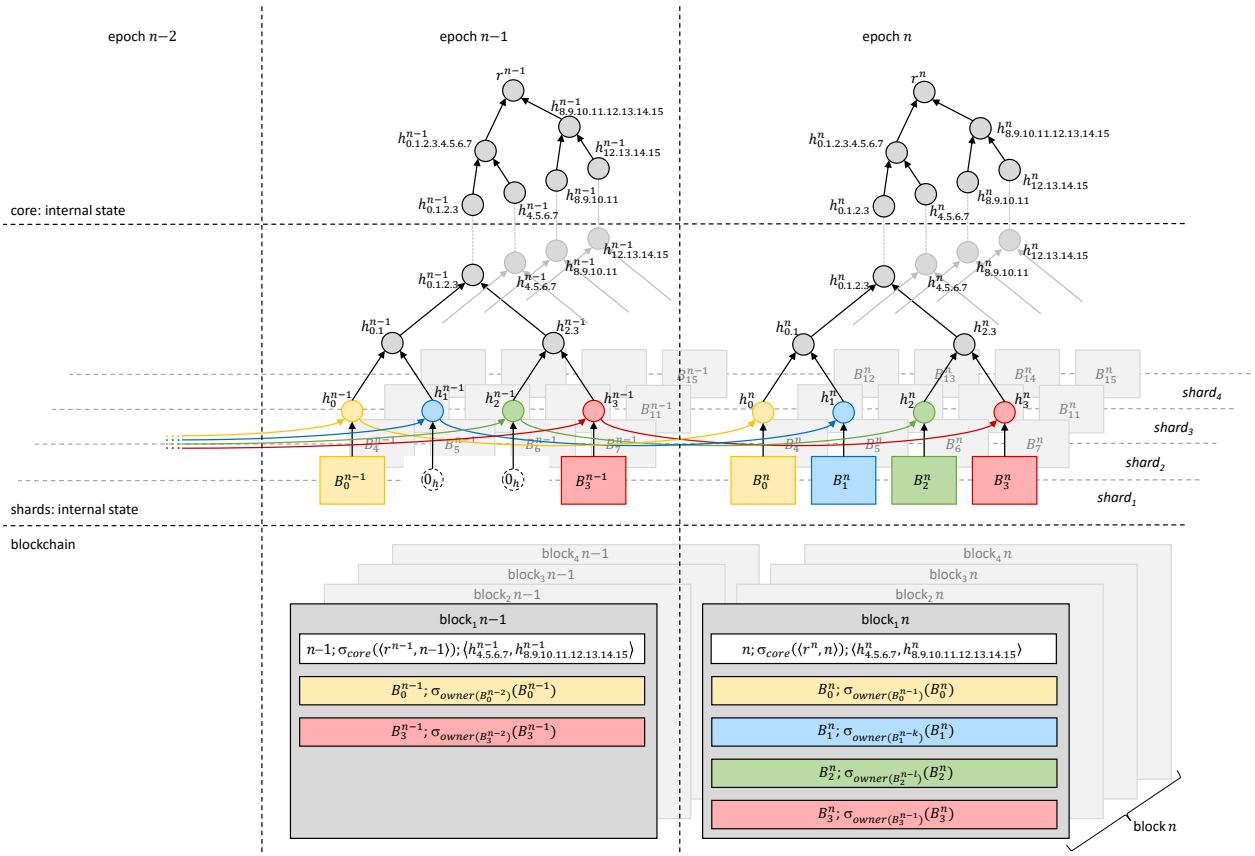

Fig. 2. KSI Cash data structures.

where $owner(B_i^k)$ is the public key address of the new owner and $value(B_i^k)$ the (the potentially new) value of the bill. Therefore, you can think of a payment order B_i^k, that is integrated into the ledger, simply as the status of the bill in the respective round. We therefore also talk about a payment order B_i^k simply as the bill B_i in round k.

Each payment order is signed by the most recent owner of the affected bill. We call a payment order together with a corresponding signature a *signed payment order*. Signed payment orders are the transactions that are recorded by KSI Cash blocks. In Fig. 2, we use $\sigma_o(x)$ to denote the signature of data x created by the private key belonging to the public key address o. In a given round, only a small subset of all bills is affected by payment orders. Therefore, it is necessary, to retrieve the current owner of a bill in the most recent block that contains a payment order with respect to this bill. For example, in Fig. 2, the most recent owners of bills B_0 and B_3 in round n, can be retrieved from the previous round $n-1$ as $owner(B_0^{n-1})$ and $owner(B_3^{n-1})$, whereas the owners of bills B_1 and B_2 need to be retrieved from the previous rounds $n-k$ and $n-l$ respectively, with both $k>1$ and $l>1$.

Internal State Each block needs to obtain further authentication data. We have designed a bill-based data structure that is optimized for performance, as will be analyzed in Sect. 4.3. For this, in each round, a Merkle tree is computed, which has one leaf for each bill, called the *hash of the bill* (in that round). If, in a given round, there exists a validated payment order for a bill, the corresponding leaf of the Merkle tree is computed from this payment order and the hash of the bill in the previous round, as depicted in Fig. 2 by various different colors (B_0:yellow; B_1: blue; B_2: green; B_3: red). If there exists no validated payment for a bill in a

given round, we take the hash of the bill in the previous round again as hash of the bill in the current round (shown in Fig. 2 as computation from the *zero hash*, denoted by 0_h, and the hash of the bill in the previous round). This means that data from several rounds are related to each other ("chained together") at the level of hash tree leaves (and not, as, e.g., in Bitcoin, at root hash level).

In each round, the Merkle tree, which we also call the *hash tree* of the round, is computed by the shards and a further so-called *core system*, of just *core* for short. Each shard validates payment orders and computes a Merkle tree for those bills that it is responsible for. The core then further integrates all of the shard trees into one single Merkle tree, see with Fig. 2. In Fig. 2 the hash tree of a round k consists of nodes of the form h^k_s, where s is a sequence $i_1.i_2.\cdots.i_n$ of bill identifiers that uniquely determines the node's position in the hash tree.

We use h^k to denote the hash tree of round k and r^k to denote its *root hash*. We call the root hash of the hash tree, that is computed by $shard_i$ in round k, the *root hash of* $shard_i$ (in round k). We call the upper part of the hash tree of round k, that is maintained by the core system, the *core hash tree* of round k. Note, that the root hash of a $shard_i$ belongs also, as a leaf, to the *core hash tree* of the respective round. For example, in Fig. 2, the value $h^k_{0.1.2.3}$ belongs both to the hash tree of $shard_1$ and to the *core hash tree* of round k.

Blocks and Bill Proofs The KSI Cash blocks are sharded. In a given round, each $shard_i$ builds a *(shard) block* denoted by $block_i\, k$. The block of round k, which is denoted as $block\, k$, consists of all shard blocks $block_i\, k$. In each round, the core system signs the root hash together with its round number $\langle r^k, k\rangle$ using the private key of the central authority, i.e., the respective KSI Cash system owner. We call such a signature a *root hash signature*, denoted by $\sigma_{core}(\langle r^k, k\rangle)$ in Fig. 2. As authentication data, each shard block $block_i\, k$ contains (i) the round number, (ii) the round's root hash signature and (iii) the *authentication path* [77,78,79] belonging to the shard's root hash in the core hash tree. For example, in Fig. 2, the shard root hash of $block_1\, n$ is $h^n_{0.1.2.3}$ and the authentication path of this shard root hash is the sequence $\langle h^n_{4.5.6.7}, h^n_{8.9.10.11.12.13.14.15}\rangle$.

After a payment order has been validated and successfully integrated into the block, the shard sends a respective proof to the respective wallet; either upon request of the wallet (*pull*) or automatically (*push*). We call proofs of successful payment orders also just *bill proofs*. With the given block design, each shard has the necessary authentication data to compile its *bill proofs*. A bill proof consists of (i) the bill number, (ii) the round number, (iii) the round's root signature, (iv) the hash of the bill in the previous round and (v) the complete authentication path of the bill in the current round's hash tree. For example, the bill proof of payment order B^n_0 in Fig. 2 is:

$$\left\langle 0, n, \sigma_{core}(\langle r^n, n\rangle), h^{n-1}_0, \langle h^n_1, h^n_{2.3}, h^n_{4.\cdots 7}, h^n_{8.\cdots 15}\rangle\right\rangle \tag{2}$$

The data contained in a KSI Cash bill proof is sufficient to authenticate a successful payment order, i.e., the wallet does not need to have all data of the full block or even the whole blockchain.

On Core Sharding The core system could also be sharded. In Fig. 2, the core hash tree is not sharded. However, a further tier of shards (or even tiers of shards) in the realm of the control system could be introduced, following the same pattern as exposed by the hash tree and block data structures in Fig. 2. However, core sharding would only become necessary, if the number of shards gets critically large. In balancing out concrete system designs, the number of bills per shard will be very large as compared to the number of shards. The number of bills per shard is determined by the expected performance of the system in terms of transaction throughput. Even for very large transaction loads, the number of shards can be expected to be relatively small (in regards of the need for core sharding), compare with Sects. 4.3 and 5.

4.2 The KSI Cash Exchange Service

Exchange Service and Exchange Strategies The KSI Cash data structures described in Sect. 4.1 are fully *parallelizable*. In a pure bill scheme, payment orders only change the owners of bills. Therefore, the set of all bills can be distributed over shards that can process payment orders in parallel without any

communication and coordination. If a wallet wants to enact a payment, the wallet breaks down the payment to a set of payment orders. We need to assume that the wallet owns a reasonable mix of bills with different kinds of values, in particular, small enough values, so that it can map a concrete payment onto payment orders, possibly in different shards. In practice, it will happen that some payments cannot be mapped onto a mix of payments orders, given a certain mix of bills in the wallet.

In some of the KSI Cash runs, we simply ensure via the system initialization that all wallets have enough small-value bills to run all payments. This is possible in the test environment, as we can steer which payments are sent to wallets. Furthermore, we have implemented an *exchange service* in KSI Cash. This exchange service is implemented as a *specialized wallet* that is (i) owned by the central authority (central bank), and is (ii) trusted by all of the other wallets. We call the specialized wallet of the exchange service the *exchange wallet*, and we call all other wallets *user wallets*, or just wallets, if clear from the context.

The exchange service needs to be initialized with sufficiently many small-value bills in the beginning of a test run to enable all payments between user wallets. There are two wallet *exchange strategies* with respect to the usage of the exchange service, i.e.:

- *Reactive Exchange.* Whenever a wallet cannot turn a payment into a set of payment orders, due to lack of small-valued bills, it first uses the exchange service before the actual payment.
- *Pro-Active Exchange.* After each successful payment, the wallet analyzes, whether its combination of bills is able to potentially serve all possible payments (up to the total amount possessed by the wallet). If not, the wallet immediately triggers an exchange via the exchange service.

A usage of the exchange service is called a *money exchange*. A money exchange simply consists of two payments, i.e., a payment from the user wallet to the exchange wallet (containing some large-valued bills) and a *simultaneous* back payment from the exchange wallet to the user wallet (containing smaller-valued bills in the same total amount).

Transactions and Transaction Rounds The transactions in an Alphabill transaction system might be organized hierarchically, depending on the respective domain and its technical realization. In KSI Cash, we distinguish three levels of transactions, see Table 1. At the outermost level, KSI Cash transactions are simply called *transactions*. They are triggered by users in order to transfer a value from one user wallet to another. At the next level, a transaction consists of either one or two so-called *transaction rounds*. A transaction round realizes a payment between two user wallets or otherwise a money exchange. At the lowest level, the transactions are the payment orders, which change the ownership of bills. In general, several payment orders are needed during a payment to fulfill its purpose. In the KSI Cash setting, we assume that the exchange wallet is trusted by the user wallets. Therefore, both payments of an exchange (i.e., the payment to the exchange wallet and back to the user wallet) can be enacted simultaneously, in a single transaction round.

KSI Cash Transaction Hierarchy	
transaction	user-triggered payment between wallets; consists of one or two transaction rounds
transaction round	payment between user wallets; *or* money exchange
payment order	changes the owner of a bill
Further KSI Cash Transaction Terminology	
money exchange	consists of two simultaneous payments: (i) from user to exchange wallet, (ii) trusted back payment from exchange to user wallet
transaction duration	started by a wallet sending payment orders; ends, when all bill proofs have been received; includes network travel times
tx/s	transactions per second (*transaction throughput*)

Table 1. KSI Cash transaction design and terminology.

A transaction must contain at least one transaction round, i.e., a payment between user wallets. Additionally, it might contain a money exchange as a second transaction round, according to the respective wallet exchange strategy (reactive or pro-active). Henceforth, we measure KSI Cash system throughput as *transaction throughput* (*transaction rate*), i.e., in terms of the number of *transactions per second*, which we denote as tx/s.

Transaction durations are defined via the start and end of transactions as follows: a transaction starts, when the respective wallet sends payment orders to the KSI Cash backend system; a transaction ends, when the receiving wallet has received all bill proofs of the transaction. This means, that transaction durations include *network travel times*. Network travel times depend on network performance; therefore, network travel time is an impact factor on KSI Cash performance measures (throughput) that is not KSI Cash specific, i.e., independent of KSI Cash technology. From a rigorous viewpoint, network travel times should be excluded from the duration times. It was decided to keep them in, in service of having realistic performance measures from a user perspective. In our test runs, network travel times had no large impact, e.g., a typical duration of creating a payment order by the wallet and sending it to the KSI Cash input router was in the range of 30-40ms only (as, e.g., compared to the standard round time of KSI Cash, which was set to 1,000ms, see Sect. 5.1).

Beyond the Pure Bill Scheme The exchange service described in Sect. 4.2 is a a solution that is consistent with the pure bill scheme, i.e., it works without introducing splits and joins in the money transactions. In practice, it needs to be ensured that always sufficiently many small-value bills are available in the exchange wallet. An approach to keep the number of small-value bills in the exchange wallet is to introduce strategies that trigger, in some appropriate way, also changes in the inverse direction, from the exchange wallet to user wallets and back.

A different way is introducing splits and joins. In an *extended bill scheme* splitting and joining bills is allowed; however, in such a way that cross-shard transactions are avoided entirely. This means, that the mechanisms of an extended bill scheme must be designed as in-shard only. The *dust collection* solution of Alphabill Money introduced in Sect. 6.2 achieves this and, therefore, is an extended bill scheme.

If a way to split bills into smaller-valued bills is introduced, we also need to introduce mechanisms that join bills of smaller values into fewer bills of larger values. Without any joining mechanism, after some time, the system would end up in a state of too high granularity; theoretically, in a state in which all bills are of the smallest value. However, if the bill granularity of the system gets too high, the system is flooded with too much payment orders and stops running efficiently.

4.3 Analytical Performance Estimations

In this section, we conduct some analytical time and cost estimations for the bill scheme of KSI Cash as introduced in Sects. 4.1 and 4.2. We deal with the costs of building the internal state and verifying ledger certificates. We use ρ_i to denote the number of payment orders processed by shard i in a given round. For the purpose of this analysis, we do not distinguish between payment orders that are requested by wallets and payment orders that are processed by the shards, i.e., we simply assume that all payment orders are validated correctly and will eventually be integrated into the block. We use β_i to denote the number of bills maintained by a shard and ς to denote the number of shards. See Table 2 for an overview of symbols used in this section and its results.

We use T_{hash} to denote the time needed to compute the hash of a single payment order and T_{sig} to denote the time needed to verify a payment order signature. A reasonable assumption is that the time needed to verify a signature is three orders of magnitude more than the time needed to compute a payment order hash, i.e.:

$$T_{sig} \approx 10^3 \times T_{hash} \tag{3}$$

With $T_{update,i}$, we denote the time needed by a shard i to update its internal state in a round (called *update time*), i.e., to compute the new hash tree for all the payment orders to be included into the block.

ρ_i : number of payment orders, shard i
β_i : number of bills, shard i
ς : number of shards
T_{hash} : hash computation time
T_{sig} : signature verification time
$T_{update,i}$: state update time, shard i
$T_{process,i}$: round processing time, shard i
$T_{process}$: total round processing time

$$T_{sig} \approx 10^3 \times T_{hash} \tag{3}$$

$$T_{update,i} \approx \rho_i \cdot \log_2 \beta_i \cdot T_{hash} \tag{7}$$

$$T_{process,i} \approx \rho_i \cdot \left(T_{sig} + \log_2 \beta_i \cdot T_{hash}\right) \tag{8}$$

$$T_{process} \leqslant \log_2 \varsigma \cdot T_{hash} + \max_i T_{process,i} \tag{9}$$

Table 2. KSI Cash time and cost estimations: symbols and summary.

The time $T_{update,i}$ corresponds directly to the size of the computed hash tree, i.e., its number of nodes. The extreme (and unlikely) case that there are as many payment orders as bills yields the following as an upper bound for the update time (according to the full number of nodes in the hash tree):

$$T_{update,i} < 2\,\beta_i \cdot T_{hash} \tag{4}$$

If there are less payment orders than bills, accordingly less leaf nodes and possibly also (many) inner nodes of the hash tree do not have to be computed. Remember from Sect. 4.1 that the previous round's hash is taken for bills for which there is no payment order in the current round. Saving hash tree computations propagates through the tree, i.e., in updating the hash tree, the hash values from the previous round's internal state can be re-reused wherever possible. (You start with computing the hashes for existing payment orders, flag the inner nodes that get affected at the next level, proceed with those nodes in the next step, and so on, until you reach the root.) Therefore, the extreme case that there is only a single payment order yields the following lower bound for the update time (according to the length of the path from a leaf to the root in the hash tree):

$$T_{update,i} \geqslant \log_2 \beta_i \cdot T_{hash} \tag{5}$$

Now, given the number of payment orders ρ_i, we can give a better upper bound for the update time than (4) as follows:

$$T_{update,i} \leqslant \rho_i \cdot \log_2 \beta_i \cdot T_{hash} \tag{6}$$

We can assume that we have very few payment orders as compared to the number of bills in a shard. Similarly, we can assume that the payment orders are usually widely "distributed", i.e., that there is no significant overlap of the paths of successful payment order leaves in the hash tree. Altogether, we can see

that, therefore, the upper bound (6) is a actually a good estimation of the update time, in general, and therefore we state:

$$T_{update,i} \approx \rho_i \cdot \log_2 \beta_i \cdot T_{hash} \tag{7}$$

If we assume that there are much less payment orders then bills, i.e., $\rho_i << \beta_i$, we see that the number of payment orders ρ_i clearly dominates the update time (7).

As part of payment order validation, the shard has to verify the payment order signatures as provided by the most recent owners of the respective bills. For the moment, we neglect further payment order validation costs. With $T_{process,i}$ we denote the processing time of a shard i during a round, which encompasses the update time plus the signature verification times, i.e., we estimate:

$$T_{process,i} \approx \rho_i \cdot \left(T_{sig} + \log_2 \beta_i \cdot T_{hash}\right) \tag{8}$$

If we assume that T_{sig} is a thousand times more costly than T_{hash}, see (3), we have that the processing time in (8) is clearly dominated by the signature verification time T_{sig} for any realistic number of bills.

Finally, we can give an upper bound for the *total processing time* of a round as follows:

$$T_{process} \leqslant 2\,\varsigma \cdot T_{hash} + \max_i T_{process,i} \tag{9}$$

As shard computations can be parallelized, the total processing time in (9) needs to consider only the slowest shard (i.e., the one that needs the maximum processing time). In addition, the total processing time needs to incorporate the update time needed for computing the core hash tree. We can assume that all shard root hashes change during a round, i.e., that each shard receives at least one transaction. Therefore, the full size of the core hash tree is the adequate basis for estimating its update time. Anyhow, as discussed in Sect. 4.1, we can consider the size of the core hash tree as relatively small, so that the shard processing time is the dominating component in (9).

5 KSI Cash Performance Evaluation

5.1 The KSI Cash Test Implementation

Overview In order to evaluate the KSI Cash concepts and architecture, in the first place in regards of performance and scalability, but also in regards of availability, a first version of KSI Cash and a respective test bench have been implemented. Particular effort and resources have been invested into simulation of a realistic scenario. This involves a realistic wallet simulation as well as the implementation of a full-fledged backend system with professional input and output sub-systems that are needed if the system has to operate in the field under real-world conditions. For example, TLS (Transport Layer Security) was imposed throughout the test runs, and session authentication was conducted for incoming payment orders.

We conducted a series of test runs under simulation of realistic usage as described in Sect. 5.2. However, we also wanted to show that the system concept is, basically, capable of unlimited performance, by showing linear scalability of the involved core data structures and algorithms. Therefore, we also conducted a series of test runs with a cut-down version of the backend system and a simplified load test generation, simply to achieve ultra-high load with still reasonable testing costs. We describe these ultra-scalability test runs in Sect. 5.3. We start with a description of the backend system as used in the realistic usage scenario in Sect 5.1 and will explain the differences of the system as used in the ultra-scalability test runs in Sect. 5.3.

Both the KSI Cash implementation and the KSI Cash wallet simulation and load test generation have been deployed, as cloud-native applications, to the Amazon EC2 cloud[19].

The KSI Cash backend system, has been implemented mainly in Go[20]. The implementation encompasses approx. 105,000 (uncommented, non-blank) lines of code (LOCs), with 87,704 LOCs in Go, plus additional

[19] `https://aws.amazon.com/ec2/`
[20] `https://go.dev/`

code in diverse programming and DevOps-related languages, see Table 3. The KSI cash wallet simulation has been implemented mainly in JavaScript. Together with auxiliary test bench code, it encompasses approx. 7,500 LOCs, with 4,322 LOCs in JavaScript, plus additional code in diverse languages, see again Table 3. The software development team consisted of ten Guardtime software developers (all senior full-stack plus DevOps software engineers) and the development efforts summed up to about ten person years.

KSI Cash Backend Code				
Language	**Files**	$\mathbf{L}_{blank}$	$\mathbf{L}_{comment}$	$\mathbf{L}_{code}$
Go	659	14,724	14,138	87,704
CSS	1	2,452	2	8,619
YAML	88	384	406	3,571
JSON	8	0	0	2,398
Bourne Shell	35	227	152	957
HCL	17	95	23	589
Python	13	203	101	537
Protobuf	35	165	669	394
Markdown	9	113	0	322
Dockerfile	13	45	4	161
HTML	2	14	2	130
make	1	22	2	61
Bash	1	3	4	13
Sum	887	18,447	15,513	105,461
KSI Cash Wallet Simulation and Test Bench Code				
Language	**Files**	$\mathbf{L}_{blank}$	$\mathbf{L}_{comment}$	$\mathbf{L}_{code}$
JavaScript	80	591	1,919	4,322
LESS	24	213	7	1,232
Java	21	202	1	848
YAML	7	47	5	454
Markdown	5	97	0	179
Bourne Shell	1	23	36	126
Gradle	5	15	1	110
JSON	4	0	0	79
DOS Batch	1	26	2	76
make	3	19	14	42
SVG	8	0	1	42
HTML	2	9	20	39
Dockerfile	2	8	1	22
Sum	163	1,250	2,007	7,571

Table 3. Implementation of KSI Cash and its wallet simulation in terms of lines of code, i.e., uncommented, non-blank lines $\mathbf{L}_{code}$, comment lines $\mathbf{L}_{comment}$, and blank lines $\mathbf{L}_{blank}$.

The KSI Cash backend system has been deployed on a total of 1,382 virtual CPUs with 8.27 TB RAM and 26.1 TB storage, compare with Table 4; and the wallet simulation has been deployed on 826 virtual CPUs with 1.8 TB RAM and 24.5 TB storage. Some smaller amount of extra resources was needed to run auxiliary modules of the test bench, adding up to 24 virtual CPUs with 88 GB RAM and 1.8 TB storage, see Table 4.

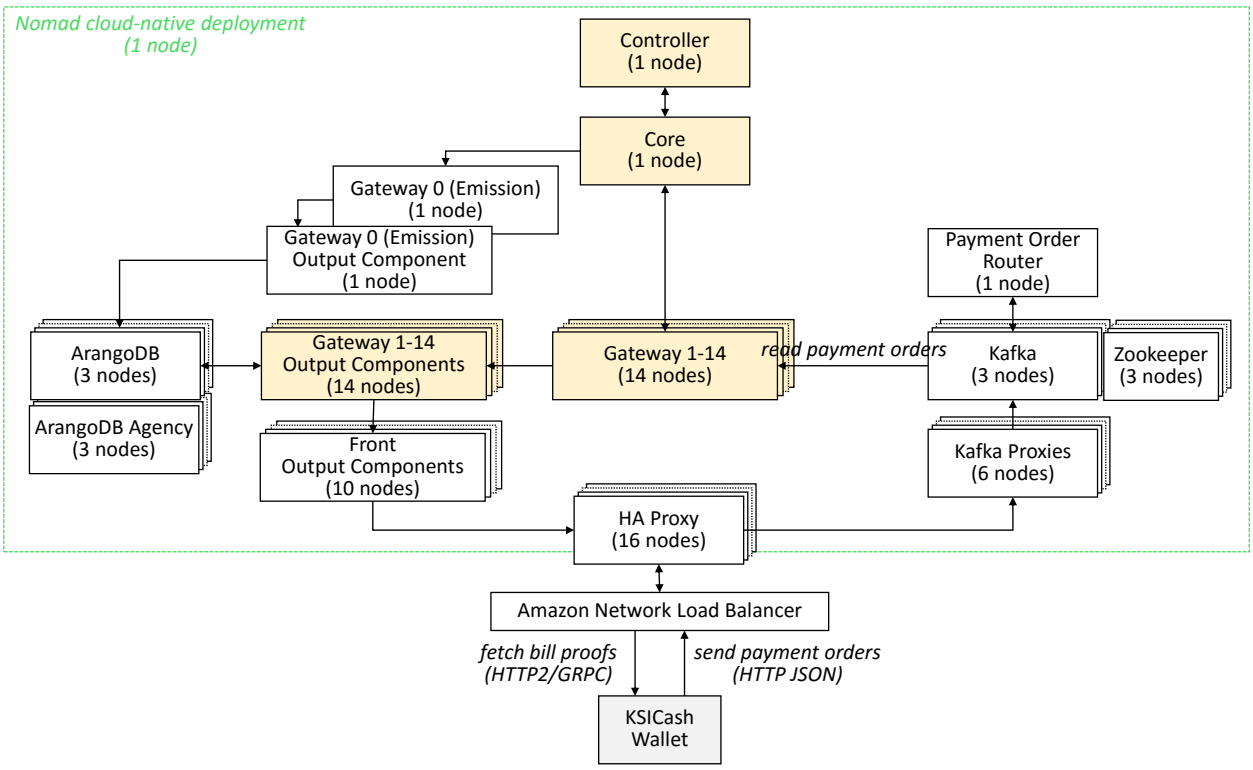

Fig. 3. KSI Cash backend system architecture.

The KSI Cash Backend System Figure 3 shows the nodes of the KSI Cash backend system as deployed in Amazon EC2. The KSI Cash wallet depicted in Fig. 3 does not belong to the KSI Cash backend system – for the architecture of the KSI Cash wallets simulation, see Fig. 4. The KSI Cash wallet sends payment order requests to and fetches bill proofs from KSI Cash through a high-availability proxy (HA proxy) consisting of 16 nodes (virtual machines). See Table 4 for technical specifications of deployment nodes, including the number of virtual CPUs (vCPUs), RAM size, storage size, processor type, Amazon EC2 instance type and the number of virtual machines (VMs). Henceforth, we use *nodes* and *virtual machines* as synonyms.

Handling payment orders in KSI Cash is sharded as described in Sect. 4. The system consists of 14 shards. According to this, there are 14 so-called *gateway* nodes (Gateway 1-14) and, additionally, 14 corresponding *gateway output components*. The gateways and their output components can be considered the workers of the system, as they are responsible for building the shard hash trees, creating blocks and issuing bill proofs as described in due course. This explains the relatively large amount of compute resources dedicated to the gateways and their output components, with 896 virtual CPUs and 7.17 TB RAM, i.e., approx. 65% of vCPUs and 87% RAM of the total compute resources of the backend system (1,382 vCPUs, 8.27 TB RAM). We name the gateways and their output components together with the *core* module and the *controller* module the *key components* of the system, as these components are responsible for all of the tasks described in Sect. 4. The remaining components allow for the initialization and management of the system and ensure its client connectivity.

Payment orders are queued in data pipelines based on Apache Kafka[21]. We use a configuration with three Kafka nodes, six Kafka proxy nodes and three Zookeeper nodes. Such configuration is standard; the three Kafka nodes are in service of availability, not because of performance optimization. There exists one queue

[21] `https://kafka.apache.org/`

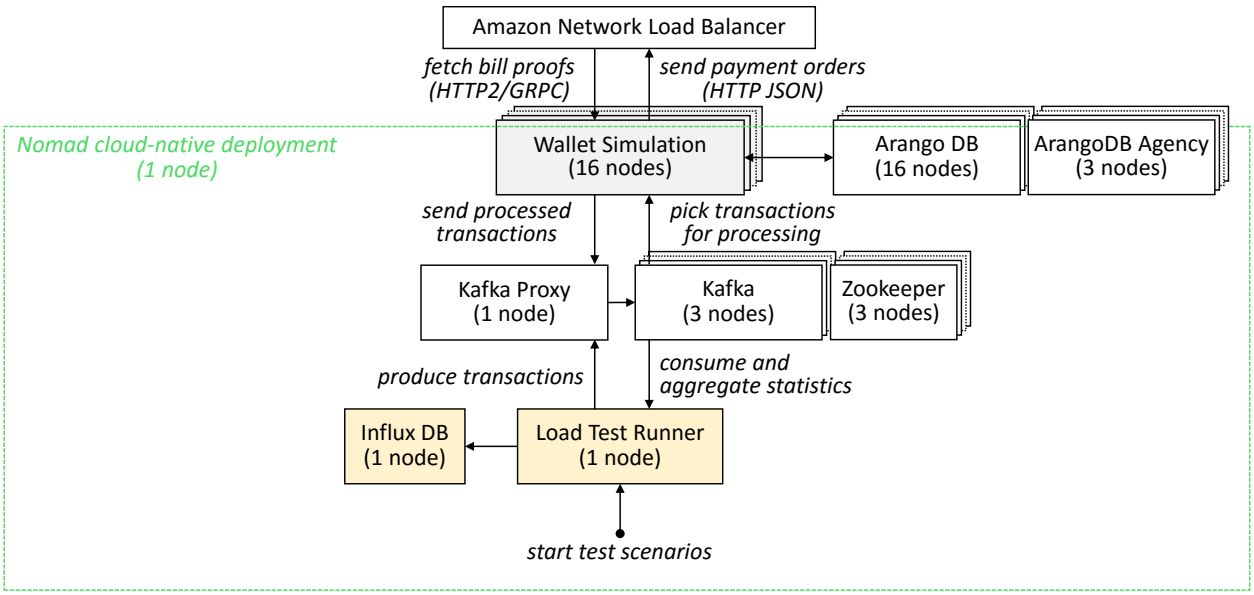

Fig. 4. KSI Cash wallet simulation and load test generation.

for each gateway in the Kafka configuration. The *payment order router* analyzes incoming payment request and routes them to the respective gateway queues. The payment order router is a potential bottleneck for high loads of payment requests, however, more parallel payment router nodes can be arbitrarily added to deal with higher loads, as payment order routing is fully parallelizable. For the loads during our test runs, one node (with 8 vCPUs) was sufficient.

In the KSI Cash implementation, gateways are synchronized, i.e., they process payment orders and build blocks in rounds of pre-determined, fixed length called *round time*, where the round time is a tuning parameter of the system. In the test runs of KSI Cash that we present in this paper, the round time was set to a length of 1,000ms.

At the beginning of each round, each gateway pulls the payments orders, that it is responsible for, from its gateway queue. It stops pulling requests after a fixed amount of time, called *payment order reading time*, set to 300ms in the test runs of this paper. Payment orders that could not be read in the pre-determined reading time, remain in the gateway queue for processing in further rounds. All gateways need to hand over calculation to the core node at least before a so-called *core offset time*, which is set to 100ms in the test runs of the paper. The round time, the payment order reading time and core offset time are all tuning parameters of the system, that need to be configured in such way that all necessary gateway calculations are finished during a round.

The core node is responsible for building the core hash tree. As soon as a gateway has received a certificate from the core node, it sends the block of the current round to its corresponding gateway output component. The gateway output stores the block in persistent storage. For storage of the blockchain, we use the multi-model database ArangoDB[22], that we run as a database cluster consisting of six nodes, i.e., three ArangoDB nodes and further three ArangoDB Agency nodes. Furthermore, the gateway output component rebuilds the hash tree of the most recent round in its internal memory, which it needs to serve bill proof requests that are sent from KSI Cash wallets. The gateway output components serve bill proof requests via a scalable tier of ten *front output component* nodes. The front output components answer bill proof requests immediately, i.e., either with the bill proof or otherwise a message which indicates that the bill proof is not yet available. In

[22] `https://www.arangodb.com/`

case a bill proof is not yet available, the respective wallet has to retry the bill proof request after a reasonable time interval, which again is a tuning parameter.

An initialization system was built for setting up the testing environment from scratch, including deployment of components, configuring a new KSI Cash instance, emitting bills, initializing and pre-funding the wallets, in whatever configuration necessary. The Gateway-0 node, also called *emission gateway* and its output component are used to run this system initialization. The whole KSI Cash backend system has been developed as a cloud-native application. We have decided to use Nomad[23] as container orchestration system for KSI Cash. One node is dedicated as Nomad server for the KSI Cash backend.

The KSI Cash Wallet Simulation Figure 4 shows the components of the KSI Cash wallet simulation. The test developer starts test scenarios via the *load test runner* component. For each test run, the load test runner generates a series of transactions among wallets. These transactions are buffered in a Kafka configuration consisting of three Kafka nodes plus one Kafka proxy node and three Zookeeper nodes. The *wallet simulation* is run by 16 nodes, connected with an ArangoDB database to hold the wallet data, run by 16 ArangoDB nodes plus 3 ArangoDB Agency nodes.

The wallet simulation has a series of tasks. First it picks transactions from the Kafka queuing system. In case of test runs including exchange services, it breaks down transactions to transaction rounds and also steers the execution of these transaction rounds. It splits transaction rounds into payment orders and sends these to the KSI Cash backend system. Furthermore, the wallet simulation requests bill proofs from the KSI Cash backend system. A simulated wallet waits a reasonable time after sending a payment order and before requesting the corresponding bill proof, which is at least the round time plus a certain *waiting time offset*. This time offset is a tuning parameter of the overall KSI Cash system and is set to 250ms in the current setting. If the KSI Cash backend system replies to a bill proof request that the bill proof is not yet available, the wallet will again wait the specified waiting time before re-issuing the bill request.

Once a simulated wallet has received all bill proofs needed to prove a transaction, it sends it back to the load test runner as a successfully *processed transaction*. The load generator aggregates the results and submits them to the InfluxDB[24] time-series database. The test engineer can then use InfluxDB for various statistical analysis. Furthermore, InfluxDB was connected to the data analysis and visualization tool Grafana[25] for further analysis (not shown/listed in Fig. 4 and Table 4).

5.2 Realistic Usage Test Runs

Overview Several test runs have been conducted. The purpose of the *baseline test run* in Sect. 5.2 was to determine the transaction duration under very small transaction load, i.e., such that can be considered as minimally needed transaction duration. The purpose of the *production test run* in Sect. 5.2 was to set the system continuously under a high load, according to the key performance indicator of 10.000 tx/s, which was initially set for the proof of concept. The purpose of the *maximum throughput test run* in Sect. 5.2 was to push the test system to its limits, in regards of its current configuration and tuning parameter setting. The purpose of the *exchange service test runs* in Sect. 5.2 was to investigate differences of behaviors of KSI Cash's reactive versus pro-active wallet exchange strategies. We use the *peaks and lows test run* in Sect. 5.2 to investigate gateway calculation times. The purpose of the memory consumption test runs in Sect. 5.2 was to investigate the RAM usage of gateway nodes and gateway components under several different transaction loads. The *24 hours test run* in Sect. 5.2 was for the purpose of demonstrating the availability of the system.

Bill and Wallet Initialization Most of the test runs in this paper (Sects. 5.2, 5.2, 5.2, and 5.2) have been initialized with the so-called *standard initialization* as follows:

– *Bills*: 2.8 billion bills of different values have been emitted, i.e., 200 million bills per gateway.

[23] `https://www.nomadproject.io/`
[24] `https://www.influxdata.com/`
[25] `https://grafana.com/`

KSI Cash Server Nodes						
Node Name	**vCPUs**	**RAM**	**Storage**	**Processor**	**EC2 Type**	**VMs**
Controller	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Core	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Gateway-0	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Gateway-0 Output Component	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Gateway 1–14	32	256GB	20GB	Intel Xeon Scalable	r5n.8xlarge	14
Gateway 1–14 Output Component	32	256GB	20GB	Intel Xeon Scalable	r5n.8xlarge	14
Front Output Components	16	32GB	20GB	AMD EPYC 7002	c5a.4xlarge	10
Payment Order Router	8	16GB	20GB	AMD EPYC 7002	c5a.2xlarge	1
Kafka	8	64GB	20GB+4TB	AMD EPYC 7000	r5a.2xlarge	3
Kafka Proxies	16	32GB	20GB	AMD EPYC 7002	c5a.4xlarge	6
Zookeeper	2	1GB	20GB	AMD EPYC 7000	t3a.micro	3
ArangoDB	16	32GB	20GB+4TB	AMD EPYC 7002	c5a.4xlarge	3
ArangoDB Agency Node	2	4GB	20GB+20GB	AMD EPYC 7002	c5a.large	3
HAProxy	8	16GB	50GB	AMD EPYC 7002	c5a.2xlarge	16
Nomad Server	2	2GB	20GB	AMD EPYC 7000	t3a.small	3
Total	1,382	8.27TB	26.1 TB	–	–	78

KSI Cash Client Nodes						
Node Name	**vCPUs**	**RAM**	**Storage**	**Processor**	**EC2 Type**	**VMs**
WM Worker	32	64GB	20GB	AMD EPYC 7002	c5a.8xlarge	16
Kafka	8	64GB	20GB+2TB	AMD EPYC 7000	t3a.xlarge	3
Kafka Proxy	4	8GB	20GB	AMD EPYC 7002	c5a.xlarge	1
Zookeeper	2	1GB	20GB	AMD EPYC 7000	t3a.micro	3
Load Test Runner	8	16GB	20GB	AMD EPYC 7002	c5.2xlarge	1
ArangoDB	16	32GB	20GB+1TB	AMD EPYC 7002	c5a.4xlarge	16
Arango Agency Node	2	4GB	20GB+20GB	AMD EPYC 7002	c5a.large	3
InfluxDB	8	16GB	20GB+1TB	AMD EPYC 7002	c5a.2xlarge	1
Total	826	1.8TB	24.5TB	–	–	45

KSI Cash Auxiliary Nodes						
Node Name	**vCPUs**	**RAM**	**Storage**	**Processor**	**EC2 Type**	**VMs**
ELK Stack	8	64GB	20GB+1.5TB	AMD EPYC 7002	r5a.2xlarge	1
HAProxy	2	2GB	20GB	AMD EPYC 7000	t3a.small	1
Nomad Backend Server	2	2GB	20GB	AMD EPYC 7000	t3a.small	1
Monitoring	2	8GB	50GB	AMD EPYC 7000	t3a.large	1
Consul	2	2GB	20GB	AMD EPYC 7000	t3a.small	3
Jump Host	2	4GB	20GB	Intel Xeon	t2.medium	1
Docker Registry	2	2GB	100GB	AMD EPYC 7000	t3a.small	1
Total	24	88GB	1.8TB	–	–	9

Table 4. The KSI Cash performance test server configuration (with node names referring to nodes in Figs. 3 and 4).

– *Wallets*: 100 million wallets; each wallet loaded with a combination of bills, so that the first payment is always possible without usage of the exchange service.

The maximum throughput test run in Sect. 5.2, the exchange service test runs in Sect. 5.2 and the memory consumption test runs in Sect. 5.2 used different bill and wallet initializations.

Transaction Generation During transaction generation, the transaction sums were sampled randomly from log-normal distribution, with arithmetic average amount of €25 and with 1-cent precision. This pattern follows Eurozone statistics[26] of combined cash and card retail payments. On average, a generated transaction triggers six payment orders. This means also, that six bill proof requests are needed, on average, to validate a transaction.

Payer wallets were chosen sequentially. Therefore, due to the test run durations, payer wallets have not been re-used, i.e., each wallet has been used for payment at most once during a single test run. Due to this (and the fact that each wallet has been initialized with a respective combination of bills, see Sect. 5.2), all transactions in the test runs of Sects. 5.2, 5.2, 5.2 and 5.2 have been single-round transactions, i.e., the KSI Cash exchange service has never been used in these test runs. On the other hand, the *exchange service test runs* provoked a series of double-rounded transactions by re-using wallet configurations from a previous test run without re-initializing the system, as will be explained in Sect. 5.2.

Receiver wallets were chosen randomly. Again, due to the test run durations, not all of the wallets have been used during single test runs.

We define the *duration of a test run* as the time from the start of the first round till the end of the last round, i.e., the test run durations do not include times for KSI Cash system initializations.

Baseline Test Run During the baseline test run, the system was loaded with 1 tx/s for a duration of 5 minutes, i.e., with a total of 600 transactions (with the standard initialization of 2.8B bills and 100M wallets). All of the 600 transactions have been successfully finished in less than 2 seconds, see Fig. 5. As the round time of the KSI Cash backend system was set to 1,000 ms, and the waiting time offset was set to 250ms, no transaction could have been faster than 1,250ms.

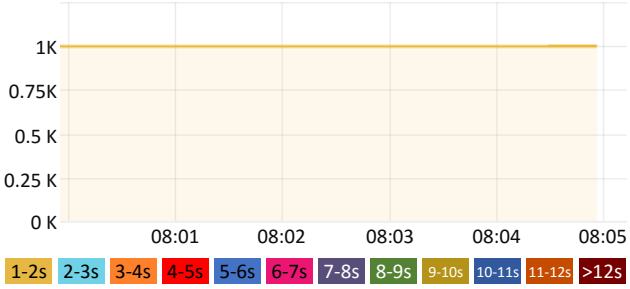

Fig. 5. Baseline test run: transactions according to transaction durations.

Production Test Run In the production test run, a steady load of 10,000 tx/s has been generated for a duration of 1 hour, see Fig. 6. The two spikes in Fig. 6 have been caused by the wallet database and can be neglected. On average, approx. 60,000 bills were used in transactions every second. Each wallet has been used at most once as the sender for a transaction. This means that a total of approx. 36 million wallets (out of the standard initialization with 100M wallets) has started transactions. This also means that the KSI Cash money exchange service has not been used and all transactions were single-round.

[26] `https://www.ecb.europa.eu/press/pr/stats/`

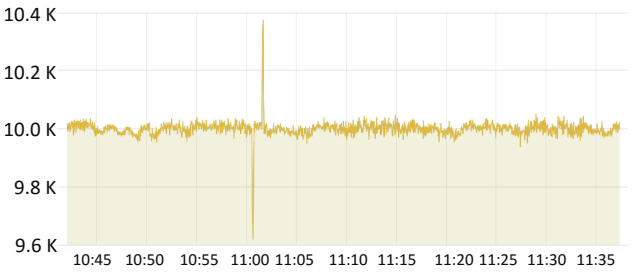

Fig. 6. Production test run: total transactions per second.

All of the transactions of this run succeeded. Figure 7 shows the distribution of transaction durations of the run: 100% of the transactions have been finished in less than 7 seconds. Actually, aside from the first spike in the run (compare with Fig. 7), all transactions have been finished in less than 4 seconds. Overall, 16.51% of the transactions have been finished in less than 2s, 95.11% have been finished in less than 3s and 99.89% have been finished in less than 4s.

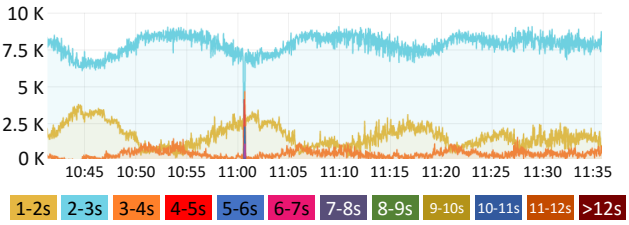

Fig. 7. Production test run: transactions according to transaction durations.

Maximum Throughput Test Run The purpose of this test run was to find the transaction rate limit of the current system setup, including the concrete setting of tuning parameters. The concept of this test run was to increase the transaction load during the test run in stages of 1,000 tx/s until the system starts failing. Here, we define *system failure* as the occurrence of a significant amount of failed transactions, i.e., on the basis of notions of *successful transaction* vs. *failed transaction*.

Definition 2 (Successful vs. Failed Transaction). *A transaction is considered a* successful transaction*, if the bill proofs of all of its payment orders have been successfully queried by a wallet during a specified* expiry time ξ*. Otherwise, it is considered a* failed *transaction.*

For the sake of this test run, we have set the expiry time for successful transactions to $\xi = 40$s (corresponding to 40 rounds of block creation in our system setting). In real-world settings, in practice, this expiry time can be surely be relaxed further; the expiry time of ξ =40sec simply shows the high ambitions of the test run. We have did not pre-specify, exactly how many failed transactions are *significant*, i.e., need to fail before we consider the situation a system failure; actually, given the success of our test run, you can assume the most strict specification, i.e., we consider the system failing upon occurrence of the first failed transaction.

A transaction rate of 15,000 tx/s of successful transactions per second has been reached. No attempt was made to re-tune the system (in terms of the KSI Cash tuning parameters) to reach even higher transaction rates with the available EC2 node deployment. It can be assumed that even higher transaction rates

could have been reached with respective tuning. Still, the result gives a stable lower bound for the possible transaction load with the EC2 node deployment described in Sect. 5.1.

The test run was initialized with 111 million bills per gateway, and 55 million wallets. The transaction load started at 10,000 tx/s. Then, the load was increased by 1,000 tx/s every 5 minutes, see Fig. 8. Up to a load of 15,000 tx/s, the system operated with a rate of 100% successful transactions.

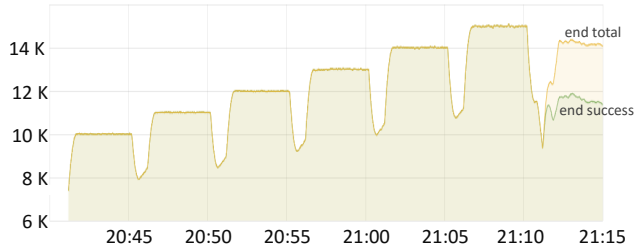

Fig. 8. Maximum throughput test run: total successful transactions per second. ('end total'=all transaction at the end of the test run; 'end success'=successful transaction at the end of the test run).

When the load was increased to 16,000 tx/s, the system started failing, see Fig. 8, where we distinguish between all transactions (*'end total'*) as opposed to successful transactions (*'end success'*) during the last five minutes of the test run. First, we can see that a significant amount of transactions failed. Second, we can see that the total number of transactions (including those failing) is less than the target load of 16,000 tx/s, and even less than the 15,000 tx/s of the previous, successful stage of the test run. This was so, because the wallet machines have been overwhelmed with requesting bill proofs at that stage, so that they failed generating the targeted number of transactions per second. Next, also the three most used gateway output components were also overwhelmed by the high number of bill proof requests at that stage. Consequentially, more and more payment orders were queued in their gateway input queues and a series of transactions failed.

We can summarize that, during system failure, the performance bottlenecks have been the wallet machines and the gateway output components. This opens opportunities for further system tuning (beyond optimizing the KSI Cash system tuning parameters) such as the introduction of a queuing layer between the front output components and the gateway output components. Again, the target of the test run was not to push the performance to the limit on the basis of the available machine, but to find a reasonable lower bound for the performance of the system with the given compute resources.

Figure 9 shows the transaction durations of this test run. Transaction durations did not increase critically with more load until the system's limit was reached. Some more transactions with a duration of up to 4s appeared with increasing load, however, even at the peak load of 15,000 tx/s, still 99.9% of transactions were still executed in less than 5s.

Exchange Service Test Runs The exchange service test runs have been conducted to test the KSI Cash money exchange mechanism and to understand its impact on the overall systems performance.

Three test runs have been conducted. The first one was a test run without usage of the exchange service, in order to have a baseline for comparison with the test runs that use the exchange service, see Sect. 5.2. The other test runs provoked some transactions that needed to use the exchange service. In the second test run in Sect. 5.2, the reactive exchange strategy was activated, compare with Sect. 4.2. In the third test run in Sect. 5.2, the pro-active exchange strategy have been activated.

Each of the three test runs was loaded with a steady transaction rate of 5,000 tx/s over a duration of 10 minutes each.

The system was initialized with 40 million bills in each gateway (a total of 560 million bills). A total number of 10 million user wallets has been created. Therefore, on average, each user wallet received 54 bills.

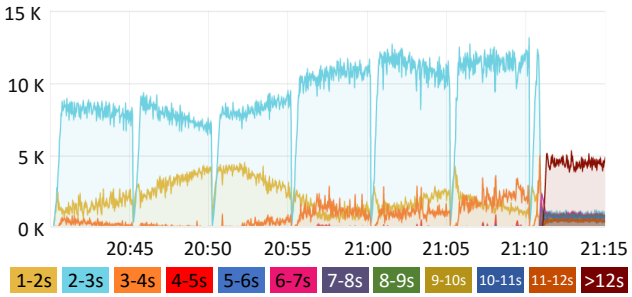

Fig. 9. Maximum throughput test run: transactions according to transaction durations.

In each wallet, the combination of bills was composed in such way that the wallet could make the first transaction without the need to use the exchange service. The exchange wallet received 20 million bills.

No Exchange Baseline Test Run During the baseline run, each sending wallet was used only once, therefore it was possible to pay all sums without the help of the exchange service. Figure 10 shows the results of the first test run. A majority of 55.07% of the transactions have been finished in 1-2s, whereas 44.87% have been finished in 2-3s no transaction took more than 4s.

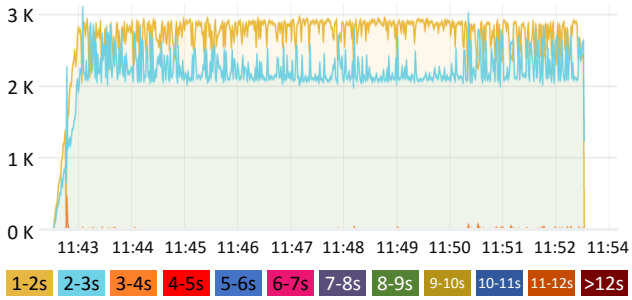

Fig. 10. No exchange baseline test run: transactions according to transaction durations.

Reactive Exchange Strategy Test Run The second test run was started after the first (baseline) test run from Sect. 5.2 without re-initialization, i.e., the wallets remained in the state as produced by the first test run. Therefore, some wallets missed some denominations for enacting a transaction at this point and had to use, reactively, the KSI Cash exchange wallet. This increased the transaction durations as can be seen in Fig. 11, when compared with Fig. 10. As opposed to the baseline test run, less transaction were finished in 1-2s (36.11%), more transaction needed 2-3s or 3-4s (42.15% and 10.3%, respectively) and transactions taking longer than 4s also occurred (>11%), see also Table 5 for a summary.

Pro-Active Exchange Strategy Test Run For the third test run, the system was re-initialized with the same data as for the first run (40M bills, 10M wallets, 20M bills in the exchange wallet, etc.). In this run, the pro-active money exchange strategy was activated, i.e., an extra money exchange round was triggered after each transaction whenever the wallet was in a state that could no longer serve directly all possible transaction requests, see Sect. 4.2. Again, compared to the baseline run in Sect. 5.2, transaction durations increased, with less transactions of 1-2s (27%), more transactions of 2-3s (68.91%) and also more transactions longer than 3s, see Fig. 12 and again Table 5 for comparison.

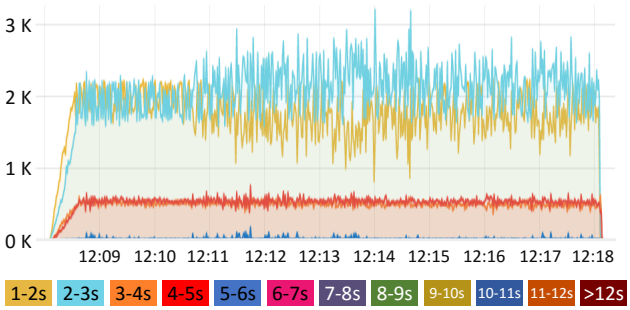

Fig. 11. Reactive exchange strategy test run: transactions according to transaction durations.

	transaction duration				
Test Run	**1-2s**	**2-3s**	**3-4s**	**4-5s**	**5-6s**
No Exchange	55.07%	44.87%	0.06%	–	–
Reactive	36.11%	42.15%	10.30%	10.94%	0.46%
Pro-Active	27.00%	68.91%	4.06%	0.03%	–

Table 5. Exchange service test runs: transaction durations of three test runs with different money exchange strategies (no exchange, reactive, pro-active).

When comparing the third test run (pro-active) with the second test run (reactive), there are much more transactions in the range of 2-3s in the pro-active run than in the reactive run. With 68.91%, it can be said that the bulk of transactions were 2-3s in the pro-active run, whereas even longer transaction occurred more often in the reactive run (>21% as opposed <5% in the pro-active run). When comparing the overall averages of the two runs, the pro-active run is slightly better (2.2712 seconds) than the reactive run (2.4735 seconds). From these results, it is not possible to decide clearly which of the money exchange strategies is better. More and longer simulations or mathematical modeling would be needed. At least, the results seem to indicate that the pro-active strategy should be preferred whenever the optimization goal is to minimize very long transaction durations (over 3 seconds).

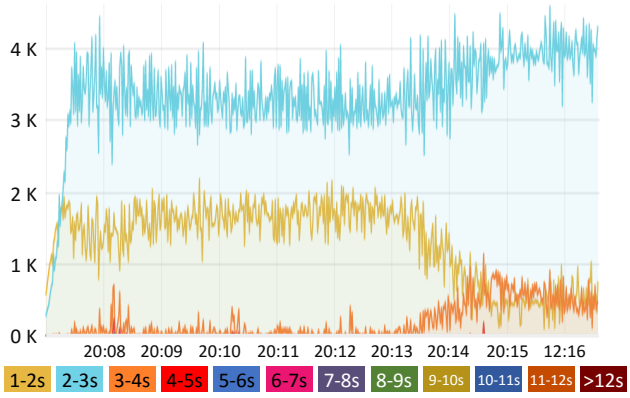

Fig. 12. Pro-active exchange strategy test run: transactions according to transaction durations.

Peaks and Lows Test Run This test consisted of six phases of different transaction rates, each with a length of 18 minutes as follows (with the system standard initialization: 2.8B bills, 100M wallets):

1. 100 tx/s
2. 1,000 tx/s
3. 10,000 tx/s
4. 100 tx/s
5. 1,000 tx/s
6. 10,000 tx/s

The idea of the test run was to simulate peaks and lows in a real world environment, in purpose of investigating gateway calculation time behavior under realistic conditions.

Figure 13 shows the result of the test run in terms of gateway *round calculation times* (for 14 gateways, one colored curve per gateway node in Fig. 13), i.e., the calculations made by a gateway during one round.

Definition 3 (Gateway Round Calculation). *A* gateway round calculation *includes:*

- *Validation of payment orders; here, processing is dominated by ECDSA (Elliptic Curve Digital Signature Algorithm) [80] signature checks. Go's cryptography library has been used for the purpose of signature checks*[27].
- *Updating the state tree of the gateway machine. This works in batch mode, i.e., some leaves of the Merkle tree are changed when the represented bills change ownership, and then the Merkle tree is re-calculated to update its root, see Sect. 4.1.*

The gateway round calculation times increase with system load. The round calculation needs to fit into the standard time window of 600ms of the test run, i.e., round time (1,000ms) minus payment order reading time (300ms) minus core offset time (100ms), see Sect. 5.1. Even under a high load of 10,000 tx/s, the gateway calculation time never exceeded 150ms (with an average of approx. 75ms), see Fig. 13.

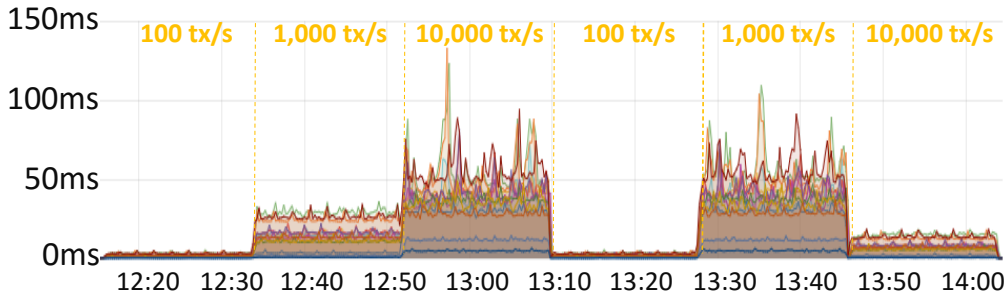

Fig. 13. Gateway round calculation time of 14 KSI Cash backend gateway nodes.

Memory Consumption Test Runs The purpose of these test runs was to investigate the memory (RAM) consumption of the gateway nodes as well as the gateway output components in regards of different amounts of bills under production condition, i.e., under 10d000 tx/s load. Three runs for have been conducted, one with 2 million bills, one with 20 million bills and one with 200 million bills per gateway, each with 100 million wallets. This means that the third run used the standard initialization (2.8B bills, 100M wallets). In each run we measured the memory usage (peak usage) of gateway nodes and gateway output components in

[27] `https://pkg.go.dev/crypto/ecdsa`

Gateway node, 8.000 payment orders per second under full load					
Nr.	**Number of Bills per Gateway**	**RAM Before Initialization**	**RAM During Initialization**	**RAM After Initialization**	**RAM Usage Under 10,000 tx/s Load**
#1	2,000,000	0.27GB	1.18	0.62GB	1.16GB
#2	20,000,000	2.89GB	12.31	6.41GB	12.42GB
#3	200,000,000	27.3GB	69.43	62.52GB	70.21GB
Gateway output component, 9.500 bill proof request per second under full load					
Nr.	**Number of Bills per Gateway**	**RAM Before Initialization**	**RAM During Initialization**	**RAM After Initialization**	**RAM Under 10.000 tx/s Load**
[1mm] #1	2,000,000	0.27GB	2.05GB	1.21GB	2.18 GB
#2	20,000,000	2.89GB	17.24GB	11.94GB	23.97GB
#3	200,000,000	27.31GB	119.43GB	117.52GB	167.07GB

Table 6. Memory consumption of gateway nodes and gateway output components.

four phases, i.e., before system initialization, during system initialization, after system initialization (when transactions have not yet been started) and during rounds with a load of 10,000 tx/s. The duration of each test run was 5 minutes (in the narrow sense, i.e., after starting the rounds).

The results of the test runs are listed in Table 6. We have arbitrarily chosen one particular, representative gateway node and its output component for Table 6. The 10,000 tx/s transaction load had a footprint of approx. 8,000 payment orders per second for the particular gateway node, and a footprint of 9,500 bill proof requests per second for the particular gateway output component. Gateway memory consumption is in positive correlation with the number of bills. The memory consumption of gateway output components is higher than memory consumption of gateway nodes. Memory consumption is higher after initialization than before system initialization. Memory consumption is higher during transactions than before the starting transactions (after initialization). Memory consumption is higher during initialization than after initialization. This is due to the garbage collector of the Go runtime environment. However, it is only slightly higher in case of many bills (test run #3). The peak memory usage of the gateway node was approx. 70GB and the peak memory usage of the gateway output component was approx. 167GB. This means that the available RAM (256GB) of the used virtual machines (compare with Table 4) has been under-utilized during these test runs.

24 Hours Test Run The 24 hours test run was conducted in service of testing availability. The system was loaded with a moderate transaction load of 575 tx/s over a period of 24 hours (with standard initialization: 2.8B bills, 100M wallets). All of the transactions have been successful, i.e., system availability during the test run was 100%. All transactions completed within 3 seconds, where the majority of the transactions completed within 2 seconds, see Fig. 14.

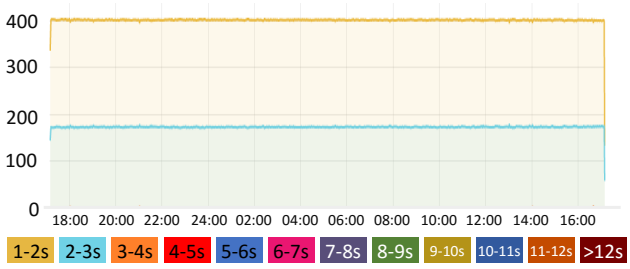

Fig. 14. 24 Hours Test Run: transactions according to transaction durations.

5.3 Ultra-Scalability Test Runs

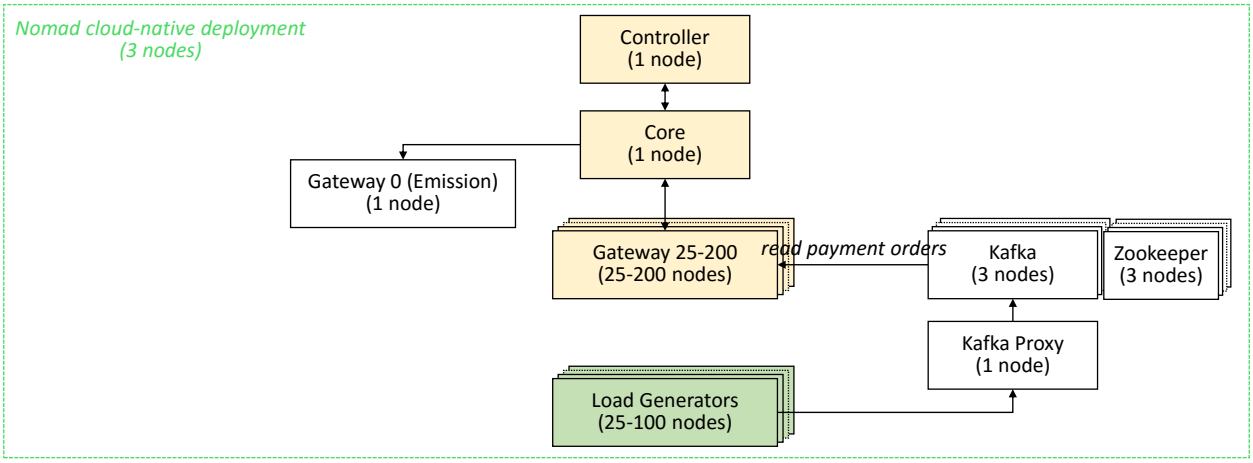

Fig. 15. The KSI cash system configuration for the ultra-scalability test runs.

Node Name	vCPUs	RAM	Storage	Processor	EC2 Type	VMs
Controller	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Core	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Gateway-0	2	4GB	20GB	AMD EPYC 7002	c5a.large	1
Gateway 1–200	8	16GB	20GB	AMD EPYC 7002	c5a.2xlarge	25-200
Kafka	8	32GB	20GB+4TB	AMD EPYC 7000	t3a.2xlarge	8
Kafka Proxy	4	16GB	20GB	AMD EPYC 7000	t3a.xlarge	1
Zookeeper	2	1GB	20GB	AMD EPYC 7000	t3a.micro	3
Nomad Server	2	2GB	20GB	AMD EPYC 7000	t3a.small	1
Load Generators	4	8GB	20GB	AMD EPYC 7002	c5a.xlarge	25-100

Table 7. Server configuration of the KSI Cash ultra-scalability performance test runs, with node names according to Fig. 15.

We have conducted a series of test runs to show that KSI Cash shows unlimited performance through linear scalability in terms of needed compute resources. We call these test runs *ultra-scalability test runs*. The target was to scale out the system to a million payment orders per second, and beyond. Conducting tests with such a high load on the full backend system and the full wallet simulation as provided by Sect. 5.1 and Sect. 5.1 would have been very costly without a clearly visible advantage of having the realistic usage simulation for this. The idea is that, together with the realistic test runs in Sect. 5.3, ultra-scalability test runs provide enough evidence for unlimited performance even if they are conducted under more simplified conditions, i.e., in what we call a *laboratory setting*.

Figure 15 shows the KSI cash system configuration for the ultra-scalability test runs, please also compare with Fig. 3. For a specification of the nodes of the system, see Table 7, compare with Table 4. A major difference is in dropping the gateway output components. This means the ultra-scalability test runs do not actually construct and store the blockchain anymore. Also, the generation of bill proofs is no longer tested. Consequentially, the ultra-scale test system also needs no front output components anymore. What is

tested, is the block creation consisting of all necessary gateway round calculations (including the validation of payment orders, and updating the state trees of the gateway machines), see Def. 3, and provision of block signatures by the core system's core component. We argue that such test design provides enough evidence for the intended scalability result.

We name the gateways together with the *core* module and the *controller* module the *central components* of the system, as together they are responsible for the block creation.

As a minor detail, the stripped-down system no longer has a high-availability proxy. Instead, the system contains load generator nodes. These are added to the system and take over the role of the KSI Cash wallet simulation from Sect. 5.1. The load generation was conceptually simplified, i.e., payment orders were generated directly, and not on behalf of (as part of) simulated transactions. The load generators were single-threaded applications run in parallel on 100 virtual machines. Each load generator was responsible for generating a load of 5,000 payment orders per second.

Some further details of the simplified conditions of the ultra-scalability test runs were as follows:

- No TLS, no session authentication for incoming payment orders; instead, the load generators send payment orders directly to the appropriate gateways.
- No bill proof generation nor bill proof requests, instead relying on the gateways' internal logging to confirm that payment orders are verified and processed. Results are collected through metrics collection.
- No storage of blocks in the ledger database. Once blocks are created by a gateway together with the associated certificates provided by the core, the blocks are just discarded.
- The gateways are relatively small, each of them handling 1 million bills.

We conducted four test runs with increasing compute resources. Table 8 shows the specification of the test runs together with the results. Tests were run with 25, 50, 100, and 200 gateways, deployed on the same number of virtual machines for each test run, with twice as many load generators for each test run. The load generators of the first three test runs have been deployed to 25, 50, and 100 virtual machines respectively, with each single load generator producing a load of 5,000 payment orders per second. The system was configured so that each gateway processed 10,000 payment orders per second. In the fourth test run, we were short of 2 million payment orders per second. This was due to the fact that we could deploy the 400 load generators to only 100 virtual machines (as in the test run with 200 load generators). Therefore, the load generators' virtual machines' CPU usage reached 100%, and they could not produce enough payment orders. Therefore, the slightly smaller number of 1,960,000 payment orders per second was not due to any KSI Cash related limitations.

Number of Gateways	**VMs (Gateways)**	**Number of Generators**	**VMs (Generators)**	**Payment orders per second**
25	25	50	25	250,000
50	50	100	50	500,000
100	100	200	100	1,000,000
200	200	400	100	1,960,000

Table 8. Linear scalability of the transaction rate (in terms of payment orders) in dependence of the number of deployed gateways.

Assuming that a payment transaction consists, on average, of six payment orders (as, e.g., in Sect. 5.2), the results mean that we have shown KSI Cash operating with more than 300,000 transactions per second.

It can be summarized that the ultra-scalability test runs demonstrate that KSI Cash is capable of processing (about) 2 million payment orders per second (corresponding to more than 300,000 transactions per second), scaling linearly in the available compute resources. The result is by no means surprising, the purpose of the test runs is merely to demonstrate that these transaction loads are possible, and beyond, even way higher transaction loads are indeed possible with still affordable compute resources.

5.4 KSI Cash Carbon Footprint

Carbon Footprint per Transaction We provide an estimation of the carbon footprint of KSI Cash on the basis of the production test run in Sect. 5.2, i.e., on the basis of 10,000 tx/s and the compute resources needed for this test run, see Table 4.

For the estimation, we consider only components related to the implementation of KSI Cash, including orchestration and edge networking, i.e. KSI Cash *server nodes* in Table 4. Load generation and analytics-related components are not included into the calculation. In total, 1,382 vCPUs (virtual CPUs) and 8.27 GB of RAM were provisioned. With respect to resource consumption, there was still room for optimization, as the weighted average CPU usage during the production test run was below 20% and the memory load was below 40%.

The experiments were run in the Amazon EC2 cloud environment. Amazon does not disclose its energy efficiency and environmental footprint in necessary detail; therefore, we use indirect calculations with the Dell EMC Enterprise Infrastructure Planning Tool[28]. For the estimation, we assemble a comparable hardware setup in the tool, choosing comparable high-end components and assuming virtualization.

In order to substitute the needed vCPU count and amount of memory, we calculate with 6 instances of Dell PowerEdge R940 servers, each configured with 4 Intel Xeon 8180 processors, each providing 28 cores and 56 threads. The servers are configured with 44 RDIMM memory modules, each 32GB, which yields 1.4 TB RAM per server. The maximal power requirement of the CPU is 205 W, and of the memory around 200 W. In total, these 6 servers amount to 1,344 CPU threads and 8.448 GB of memory. We take these 1,344 CPU threads as an equivalent of 1,344 vCPUs in our production test run, as Amazon EC2 markets one CPU thread as one vCPU.

Dell's planning tool estimates the following power consumption for our configuration:

- *in idle mode:* 304 W
- *under memory-intensive load:* 1030 W
- *under potential maximum load:* 1704 W

In our estimation, we work with 6 kW under *memory-intensive load* for our six-server configuration. We add 1 kW for network equipment and assume an equal power consumption for cooling and other data-center losses, i.e., we work with a total of 14 kW for our scenario. Power consumption for cooling and auxiliary data center resources must not be neglected [81]; they might be lower than factor *two* (depending, e.g., on the exploitation of efficient cooling techniques such a *free cooling*), but also even higher than factor *two*. Still, factor *two* is a usual rule of thumb used in practice, and we argue that it is appropriate for the purpose of our rough estimation.

Next, based on the EU's 2019 average carbon footprint of 275g CO2/kWh [29], we can conclude the following estimations for the carbon footprint of KSI Cash:

- 34t CO2 per year (with permanent load of 10.000 tx/s)[30]
- 0.0001g CO2 per transaction[31]

Our estimation does not have the ambition to predict the exact carbon footprint of the KSI Cash production system. In real-world scenarios, additional resources might be required, e.g., to ensure a higher degree of availability, or due to a larger amount of bills covering a larger money supply. Still, we suppose that, even under a pessimistic approach, this would not increase our estimate by more than 3–5 times. On the other hand, there is also the potential of energy savings, e.g., by further optimizing the KSI Cash system or using optimized hardware. To conclude, we suggest that our carbon footprint estimation is accurate enough to be used in *orders-of-magnitudes* comparisons, such as needed when comparing with Bitcoin and proof-of-work consensus, see Sect. 5.4.

[28] http://dell-ui-eipt.azurewebsites.net/

[29] `https://www.eea.europa.eu/data-and-maps/indicators/overview-of-the-electricity-production-3/assessment`

[30] $14\text{kW} \times 24\text{h} \times 365.25\text{d} \times 275\text{g CO2/kWh} = 33.749\text{t CO2}$

[31] $14\text{kWh} \times 275\text{g CO2/kWh} \,/\, 3{,}600\text{s} \,/\, 10.000\text{tx} = 0.0001069\text{g CO2/tx}$

Bitcoin Carbon Footprint The energy consumption and carbon footprint of Bitcoin and proof-of-work consensus is alarming [82]. A study [83] from the Frankfurt School Blockchain Center from 2021 estimates the carbon footprint of Bitcoin as approx. 369.5 kg CO2 per transaction. This number is based on an estimation of the annual energy consumption of 90.86 TWh (in the period from 1 eptember 2020 to 31 August 2021), an estimation of 37.97 Mt CO2 CO2 emission, and a total of 102,754,276 transactions in that period. The number of transactions is taken from the *statista.com* portal[32]. According to the *blockchain.com* portal, the number of confirmed Bitcoins transactions in that period[33] is 34,379,872, which would correspond to a carbon footprint of 1,104.4 kg CO2 per transaction (based on the 37.97 Mt CO2 annual emission estimate from above); whereas the number of confirmed Bitcoin payments is 238,745,466[34], which would amount to a carbon footprint of 159 kg CO2 per payment. Note that a Bitcoin, a transaction can contain several payments to different recipients (called *batching*). Such *batched* payments must not be confused with the payment orders of KSI Cash. In KSI Cash, a transaction consists of one or more payment orders; when comparing energy consumption of KSI Cash and Bitcoin, it is the Bitcoin's *batched payment* that corresponds to a KSI Cash transaction. Henceforth, we talk about a Bitcoin batched payment also as *payment transaction*.

A study [84] of the central bank of the Netherlands (De Nederlandsche Bank) from 2021 comes up with similar results, i.e., an annual energy consumption of 70 TWh for 2020 (45 Mt CO2; 402 kg CO2 per transaction) and 54 TWh for 2019 (36 Mt CO2; 300 kg CO2 per transaction).

Other estimates are:

- 21.5 to 53.6 Mt CO2 in 2018 [85]
- 45.8 TWh and 22 Mt CO2 in 2018 [86]
- 1,216.51 kg CO2 per transaction[35]

Similar estimates are also reported in [87].

Given these data, it can be summarized that the carbon footprint of Bitcoin is much more than 100 kg CO2 per payment transaction.

6 Essential Alphabill Platform Components

6.1 Fundamental Alphabill Design Concepts and Elements

Asset Representation Each kind of asset is implemented by an individual transaction system having individual characteristics. Each partition maintains assets *entities* (units) and enables *transactions* that create new entities, delete identities, or change the attributes of entities.

Definition 4 (Entity Representation). *Each entity is represented as a triple* $\langle \iota, D, \varphi \rangle$ *where:*

- ι *is a unique identifier of the entity.*
- D *is the data part that encompasses all relevant attributes of the entity.*
- φ *represents ownership –* φ *is a predicate (logical condition) that defines the rules and restrictions of the next transaction with the entity, i.e., so that in order to execute the next transaction* T *with the entity, the initiator has to provide an ownership proof* s *such that* $\varphi(T, s)$ *holds.*

The system identifiers are totally ordered, i.e., for each pair of identifiers ι_1, ι_2, either $\iota_1 = \iota_2$, $\iota_1 < \iota_2$, or $\iota_2 < \iota_1$.

Ownership predicates are an essential concept. For example, in Bitcoin, such ownership predicates can be found in the form of so-called locking scripts [76]. The simplest example of an ownership predicate is public-key based signature verification, i.e.,

$$\varphi(T, s) \equiv \mathsf{Verify}(\mathsf{pk}; T, s),$$

[32] `https://www.statista.com/statistics/730806/daily-number-of-bitcoin-transactions/`
[33] `https://www.blockchain.com/charts/n-transactions`
[34] `https://www.blockchain.com/charts/n-payments`
[35] `https://digiconomist.net/bitcoin-energy-consumption`

which holds if s is a signature on T that verifies with the public key pk, where the public key pk is an additional parameter of φ.

Platform Elements All kinds of transaction systems can be integrated into the Alphabill infrastructure. It is assumed that all partitions follow the Alphabill design principles as described in Sect. 3.2, in particular, achieving highest security and scalability. The Alphabill enables the integration of transaction systems and supports the adherence to its design principles by providing a series of platform elements and respective platform behavior, i.e., the Alphabill platform

- defines a *language* for describing the functionality of transaction systems (syntax and semantics of state and transactions),
- provides *libraries* and *toolkits* for developing block-chained transaction systems in the Alphabill framework,
- *registers* and *assigns identifiers* α to transaction systems, based on the descriptions of the transaction systems,
- provides *ledger certificates*, through the blockchain mechanism, i.e., irrefutable proofs of the current (block n), the parameters D of entities and the transactions T with entities.

The Alphabill development libraries and toolkits unlock the technological know-how in sharding of KSI Cash and Alphabill Money to be used for realizing massively scalable transaction systems. All transaction systems that exploit these libraries and tools will show a design consisting of shards and a core system as described in detail in Sect. 4.1, see Fig. 1.

The ledger certificates provided by the Alphabill platform are described in Sect. 6.1.

In order to integrate a transaction system as a partition into the Alphabill platform, the steps described in Sect. 6.1 have to be followed.

Ledger Certificates The Alphabill platform supports three kinds of ledger certificates:

- *Transaction Certificate.* A transaction certificate proves that a transaction T was included into the n-th block of a transaction system.
- *Data Certificate.* A data certificate proves the value of D (its attributes) of an entity ι in a block n.
- *Non-existence certificate.* A non-existence proves that the entity with identifier ι does not exist in round n. Non-existence certificates play a crucial role in atomic multi-asset swap transactions.

A ledger certificate consists of

- the *authentication path* of the state tree of the n-th block,
- the *uniqueness certificate* of the n-th block which depends on the root of the state tree and the consensus mechanism of the partition.

Alphabill partitions may be based on different kinds of state trees such as:

- *Pure Merkle tree* [77] – supports transaction certificates and data certificates, but does not support non-existence certificates.
- *Authenticated search tree* [88] – provides all three kinds of ledger certificates.
- *Count-certified hash tree* [89] – provides all three kinds of ledger certificates, and, in addition, also a proof of a *summary value* or an *invariant* of the transaction system (e.g., the total amount of money in case the transaction system represents a money scheme).

Transaction Orders In the Alphabill platform, a unified message format is used for all transaction orders.

Definition 5 (Transaction Order Format). *Each transaction order P in a transaction system is a tuple $\langle \alpha, \tau, \iota, A, T_0 \rangle$, consisting of:*

- α – *the system identifier*
- τ – *the message type identifier*
- ι – *the entity identifier*
- A – *a list of transaction attributes*
- T_0 – *a timeout specification*

As usual, we use *dot notation* to denote the several items of a tuple, e.g., we use $P.\alpha$ to denote the system identifier of a transaction order P.

The concrete format of the list of attributes A depends on the message type τ. The timeout T_0 represents the expiry time of the message in terms of block numbers, i.e., P can be accepted in block n only if $n < T_0$. Timeout specifications allow for providing reliable evidence that a transaction has not and will not be accepted by the system.

A *signed transaction order* is a pair $\langle P, s \rangle$, where P is a transaction order and s is an ownership proof, such that $\varphi(P, s) = 1$, where $\langle \iota, D, \varphi \rangle$ is an entity.

Sharding Scheme The Alphabill platform suggests a unified sharding strategy for all transaction systems. Each shard is assumed to manage a part of the partition's entities within a range of identifiers ι (see Def. 4) from an interval $[\iota_{\min}, \iota_{\max}]$.

Every shard has a shard identifier γ that is a finite bit string. Initially, the partition might start as consisting of only one single shard with the empty bit string as identifier γ. Upon a certain event trigger (e.g., a threshold number of transactions per shard), the shard with identifier γ is split into two shards with identifiers γ_0 and γ_1. The state tree of the shard is split by applying a tree-splitting algorithm. Hence, the set Γ of all shard identifiers is always a prefix-free code.

Definition 6 (Sharding Scheme). *A sharding scheme is represented as a pair $\langle \Gamma, \rho \rangle$, where ρ is a function on Γ that, given a shard number γ as input, outputs the interval $[\iota^{\gamma}_{\min}, \iota^{\gamma}_{\max}]$ of the shard γ.*

Transaction System Specification Each partition in the Alphabill network is completely defined by a set of parameters, called *transaction system specification.*

Definition 7 (Transaction System Specification). *A transaction system specification consists of*

- α – *the system identifier*
- *the type specification (meta data) of the data part D of the parition's entities*
- *a state tree type*
- *the initial sharding scheme and the shard-splitting trigger*
- Σ_0 – *the initial state of the system*
- *a set $\mathbb{T}$ of message types*
- *for each message type $\tau \in \mathbb{T}$:*
  - A_τ – *a list of attributes forming the message content*
  - *a specification of the message's effect onto the system state*

Message types could also be called *transaction message types* or just *transaction types* for short, however, *message type* is common terminology in transaction system and data interchange technologies and standards such as in the field of EDI (electronic data interchange) [90]. Given a transaction of message type τ, we call this transaction also just a τ *transaction* for short.

Joining the Alphabill Platform Joining a new transaction system into the Alphabill platform requires the following steps:

1. Describing the new system by defining the parameters of the transaction system, see Def. 7
2. Applying Alphabill toolkits to compile the system code
3. Preparing machines for each shard and deploying the code to machines
4. Registering the system and obtaining the system identifier α

The concrete registration procedure depends on design choices. In case of automatic registration, the registration is regulated by ledger rules, i.e., registration is a special blockchain transaction.

6.2 Alphabill Money

Pure bill-type money schemes only use *ownership transfer transactions*, i.e., such transactions that change only the ownership conditions of bills. Pure bill schemes enable massively parallel decomposition of the money system [20,19], but also have shortcomings. Similar to physical cash, it is not always possible for a party to pay exact amounts and therefore, some additional services, such as a money exchange service, are needed, compare with Sect. 4.2.

To realize a digital currency based on a pure bill-type money scheme, one has to start with a large number of bills existing already from day one. For a central bank money, such a solution is sufficient as a central bank (central authority) can always change the distribution of the monetary denominations by its authority. However, in the permissionless case, having too many bills might quickly become too costly for the validators of the blockchain. Here, it would be reasonable to start from a few high-value bills, enable the users to split the bills if needed and to provide an efficient mechanism to subsequently join bills. This idea is the design rationale behind our proposed *dust collection* solution in Sect. 6.2 and 6.2.

Split Payments An *extended bill money scheme* addresses the shortcomings of pure bill schemes by introducing *split* payments that make sure that exact payments are always possible. When applied to a bill, a split transaction reduces its value by an amount n and creates a new bill with that value n.

Split type transactions enable exact payments but introduce the critical problem that too many small-value bills (*dust bills*) emerge over time. Therefore, additional transactions and ledger mechanisms are needed to reduce the amount of dust bills by joining them to larger bills. We address the issue of dust bills in Sect. 6.2 by introducing new types of transactions along with a specialized type of bill.

Dust Collector and Bill Swap Payments In the extended bill scheme of Alphabill Money, we introduce a special type of ownership – the *dust collector* (DC). Bills owned by the DC are not in the usual money circulation but can be considered as frozen by the system. For the purpose of consolidating dust bills, further transactions with the DC-owned bills are conducted automatically by the ledger operation rules (block creation rules) and, hence, DC represents a built-in smart contract (a smart contract built in the system).

Furthermore, there is a special DC-owned bill with identifier ι_0 called *dust collector money supply*. The system is sharded and every shard must have its own DC money supply.

Users who want to get rid of their dust bills, can transfer them to DC via special ownership transfer transactions of message type transDC and get ledger certificates of them. These transactions are processed by the system in such a way that a new bill with value n is issued to the owner specified in the transaction orders, and simultaneously, the dust collector money supply (the value of the bill ι_0) is reduced by n.

By presenting those certificates to the system, users can then obtain new larger bills via bill swap transactions. Each bill swap transaction contains a list of transDC transactions with ledger certificates and an owner condition φ. Once a bill swap transaction is received by a shard of the Alphabill native currency partition, the ledger certificates of the listed transDC transactions are verified and a new bill is created with the sum value of the dust bills and with φ as owner. A bill swap transaction has a unique identifier that uniquely

points to the shard in which it can be executed. This prevents that the same bill swap transaction is used in two different shards, i.e., double-spending the bill swap transaction.

Transfers to DC and corresponding swaps alone do not reduce the number of small-value bills in the system. In order to achieve a consolidation of small-value bills, we need to introduce a mechanism for joining dust bills, as we will discuss in due course in Sect. 6.2.

Bill Swap Scenario To enact a DC swap, a user wallet selects a set of dust bills with identifiers $\iota_1, \ldots, \iota_m$ and values $v_1, \ldots, v_m$, respectively. The bill swap procedure consists of the following steps:

1. The user wallet computes
$$\iota \leftarrow h(\iota_1, \ldots, \iota_m),$$
where h is a cryptographic hash function.
2. For every dust bill ι_k, the wallet creates a signed $\mathsf{transDC}$ transaction order P_k that contains ι, an ownership condition a for a potential exchange bill, a timeout T_0, and sends P_k to the system.
3. It might be that some of the payment orders do not reach the system in time. Let $P_{k_1}, \ldots, P_{k_\ell}$ be the payment orders that were successfully received by the system. For these payment orders, the wallet obtains ledger certificates (transaction certificates) $\Pi_{k_1}, \ldots, \Pi_{k_\ell}$.
4. For the bills (from the list $\iota_1, \ldots, \iota_m$) that have not reached the system in time, the wallet obtains ledger certificates (data certificates) that the ownership in block $n \geqslant T_0$ has not changed. This ensures that these payments indeed failed.
5. The wallet sends a signed bill swap transaction $\langle P, s \rangle$ to the system, where
$$P = \langle \alpha, \mathsf{swap}, \iota, A, T_0 \rangle,$$
$$A = \langle a, \iota_1, \ldots, \iota_m; P_{k_1}, \ldots, P_{k_\ell}, \Pi_{k_1}, \ldots, \Pi_{k_\ell}, v' \rangle,$$
and
   - a is the ownership of the potential new exchange bill,
   - $\iota_1, \ldots, \iota_m$ are the identifiers of the dust bills,
   - $P_{k_1}, \ldots, P_{k_\ell}$ are the successful payment orders of type $\mathsf{transDC}$,
   - $\Pi_{k_1}, \ldots, \Pi_{k_\ell}$ are the ledger certificates of $P_{k_1}, \ldots, P_{k_\ell}$,
   - $v' = v_{k_1} + \ldots + v_{k_\ell}$ is the value of the exchange bill.
6. The system checks the following required conditions:
   (i) $v' = v_{k_1} + \ldots + v_{k_\ell}$, where $v_{k_1}, \ldots, v_{k_\ell}$ are the values of the corresponding bills (v_{k_i} is included in P_{k_i}).
   (ii) $v' \leqslant v_0$, where v_0 is the current value of the DC money supply, i.e., where there exists a sufficient DC money supply.
   (iii) there exists no bill with identifier ι. If the hash function h is collision-resistant, this condition only fails in case of an attempt to use the same swap message twice, possibly, to maliciously double spend it.
   (iv) $\{P_{k_1}.\iota, \ldots, P_{k_\ell}.\iota\} \subseteq \{\iota_1, \ldots, \iota_m\}$ – all bill identifiers in payment orders are elements of $\{\iota_1, \ldots, \iota_m\}$
   (v) $\iota = h(\iota_1, \ldots, \iota_m)$ – the identifier ι of the new bill is properly computed.
   (vi) $P_{k_1}.\tau = \ldots = P_{k_\ell}.\tau = \mathsf{transDC}$ – bills were transferred to DC
   (vii) $P_{k_1}, \ldots, P_{k_\ell}$ all contain ι.
   (viii) the payment orders $P_{k_1}, \ldots, P_{k_\ell}$ all contain the proper ownership condition a (of the potential new exchange bill),
   (ix) $a(P, s) = 1$, i.e., s is the signature of a on P,
   (x) $\Pi_{k_1}, \ldots, \Pi_{k_\ell}$ are proper ledger certificates for $P_{k_1}, \ldots, P_{k_\ell}$.
7. If all of the required conditions have been verified, a new bill, represented by the triple $\langle \iota, v', a \rangle$, is created.

Now, let us analyze the security of the bill swap procedure.

It is not possible to obtain two different exchange bills (say, with identifiers $\iota \neq \iota'$) for the same set of dust bills, because the identifier ι is a deterministic function of $\iota_1, \ldots, \iota_m$.

The same swap cannot be used twice in a shard, because the swap transaction with bill identifier ι is not accepted whenever a bill with the identifier ι already exists.

The same bill swap transaction cannot be used in two different shards, because the sharding scheme uniquely defines the shard of ι.

The ownership a of the new bill cannot be modified by outsiders, because a is included in the payment orders $P_{k_1}, \ldots, P_{k_\ell}$ and this inclusion is verified.

It is not possible that swap dust bills are controlled by somebody else, because the signature s of the swap order has to contain a signature s that satisfies a.

If some (or even all) of the transDC payment orders are not received in time, their ownership does not change and they can be swapped next time. If just one transDC payment order goes through, the owner of this dust bill just receives an equivalent dust bill in a (possibly) different shard. Hence, there is no risk of losing money during a swap.

Dust Collection Process In the extended bill scheme, dust collection is introduced as a necessary automatic functionality related to block creation, i.e., regularly (in a way well-defined by the ledger rules), each block creator has to delete a set

$$(\iota_1, v_1, DC), \ldots, (\iota_k, v_k, DC)$$

of dust bills and, simultaneously, raise the DC money supply (ι_0, v_0, DC) accordingly by increasing its value by

$$d = v_1 + \ldots + v_k.$$

All the activities related to dust collection preserve the total money of the system, including DC money.

The DC money is a purely technical, system-related measure and does not actively participate in the usual "business transactions", i.e., it is not meant to be directly accessed by system users.

The State Tree of Alphabill Money Alphabill Money uses an AVL-type authenticated search tree, see Adelson-Velsky and Landis [91], where the indices are provided by the bill identifiers. In addition, it is a count-certified tree where the counters in nodes represent the total value of the corresponding sub tree. All nodes, not only the leaf nodes, contain data about bills.

The tree is represented as a pair $\langle \iota_r, N \rangle$, where ι_r is the bill identifier of the root node and N is an indexed array (dictionary) of nodes. We use $N[\iota]$ to denote the node that is associated with the bill with identifier ι. If there is no bill with identifier ι, we write $N[\iota] = \bot$.

Definition 8 (Alphabill Money State Tree Node). *An Alphabill Money state tree node $N[\iota]$ is a tuple $\langle \varphi, v, x, V, h, \iota_L, \iota_R, d, b \rangle$ consisting of the following:*

- *φ – bearer condition*
- *v – value of the bill*
- *x – bill hash computed as $H(x', T)$, where H is a hash function, T is the last transaction order with the bill ι and x' is the previous bill hash of the same bill.*
- *V – summary value computed as*

  $$V \leftarrow v + N[\iota_L].V + N[\iota_R].V,$$

  where we assume that $N[\bot].V = 0$
- *h – hash of the node, computed as*

  $$h \leftarrow H(\iota, H(\varphi, v, x), V; h_L, V_L; h_R, V_R),$$

  where h_L, h_R are the hashes of nodes $N[\iota_L], N[\iota_R]$, and V_L, V_R are the summary values of $N[\iota_L], N[\iota_R]$ under the assumption that the hash of $\bot$ is the zero-hash 0_h
- *ι_L – left node identifier (can be $\bot$ if there is no left node)*
- *ι_R – right node identifier (can be $\bot$ if there is no right node)*

– d – *depth of the sub-tree computed by*

$$d \leftarrow \max\{N[\iota_L].d, N[\iota_R].d\} + 1,$$

assuming that $N[\bot].d = 0$.

– b – balance factor *[91] computed by*

$$b \leftarrow d_L - d_R,$$

where d_L *and* d_R *are the depths of* $N[\iota_L]$ *and* $N[\iota_R]$*, respectively, assuming that* $N[\bot].d = 0$.

In Def. 8, the value of $N[\iota_r].V$ represents the total amount of money in the system (or shard) and $N[\iota_r].h$ is the root hash of the system (or shard).

6.3 The Alphabill Atomicity Partition

The Goal of Atomicity Control Let $u_1, \ldots, u_m$ be entities with identifiers $\iota_1, \ldots, \iota_m$ and owner conditions $\varphi_1, \ldots, \varphi_m$, respectively. The entities $u_1, \ldots, u_m$ may belong to different transaction systems (partitions) with identifiers $\alpha_1, \ldots, \alpha_m$, respectively. It is assumed that in each of these partitions, there is a kind of transaction available for changing the ownership conditions of entities.

The goal is to transfer the entities to new owner conditions $\varphi'_1, \ldots, \varphi'_m$ *atomically*, i.e., in a way that either

– *all of its component transfers take place* – all entities $u_1, \ldots, u_m$ are transferred to the new owner conditions $\varphi'_1, \ldots, \varphi'_m$, or
– *none of its component transfers take place* – all entities will have owner conditions equivalent to the previous conditions $\varphi_1, \ldots, \varphi_m$.

All entities may potentially be controlled by different parties. We assume that these parties may communicate in order to agree on the atomic transfer, i.e., after communication, all parties know all of

$$\iota_1, \ldots, \iota_m, \varphi_1, \ldots, \varphi_m, \alpha_1, \ldots, \alpha_m.$$

The parties also agree on other transaction specific parameters.

If there is more than one party involved, such transfer is an *atomic swap*. If there is only single party involved, it is an *atomic multi-entity, single-asset transfer*.

Implementation Restrictions To implement such a transaction, the parties send *component transaction orders* $T_1, \ldots, T_m$ with the new owner conditions $\varphi'_1, \ldots, \varphi'_m$ to the system. These new conditions have to be designed in a way that the atomicity condition is satisfied, i.e., the new owner of every entity ι_i can execute the next payment only if there is evidence that all other transactions $T_1, \ldots, T_{i-1}, T_{i+1}, \ldots, T_m$ have been accepted and included in the ledger. The previous owner of ι_i can execute the next transaction only if there is evidence that at least one other transaction has not been accepted.

In distributed databases technology, such orchestration of multi-shard transactions is usually achieved with two-phase commit protocols. In a partitioned blockchain technology, two-phase commits have to be implemented in a certified way – before executing a component transaction in one partition the status of other component transactions in other partitions have to be certified, i.e., there has to be evidence about the status of other component transactions.

Basically, there are two possible approaches. A usual approach used in known blockchain technologies is to define the new owner conditions $\varphi'_1, \ldots, \varphi'_m$ (locking scripts) in a way that every φ'_i includes direct verification of the status of all other transactions $T_1, \ldots, T_{i-1}, T_{i+1}, \ldots, T_m$. A shortcoming of this approach is that the conditions φ'_i are large in size and complex to verify.

Therefore, we use a different approach, where the status of the multi-component transaction is managed by a special stateful smart contract called *atomicity agent*. Instead of using direct verification of status of

other component transactions, the condition φ'_i verifies the status of the multi-component transaction via checking the status of a particular atomicity agent.

The information about the status must be consistent, i.e., always the same for all verifying parties. Therefore, a single server implementation of atomicity agents is insufficient, because a malicious server may give inconsistent status information to the parties, resulting in a violation of the atomicity condition. Implementing the atomicity agent as a smart contract in a blockchain prevents inconsistent status information because of the uniqueness property of blockchains. Therefore, all atomicity agents are managed in a special partition of the system – the *atomicity partition*.

To implement this kind of swap transactions, we contribute a *three phase commit protocol*, described in Sect. 6.3.

Atomicity Partition Description There is a specific transaction system (partition) with identifier α_0 that provides necessary unique references for atomic multi-asset swap transactions. We call it the *atomicity partition*.

Entities of the atomicity partition are the atomic multi-asset swap transactions, i.e., each of these transactions has a unique pseudo-random identifier ι (referred to as *contract identifier*) in the atomicity partition. Furthermore, each entity has a data part D containing:

- a list $\langle \alpha_1, \iota_1 \rangle, \ldots, \langle \alpha_m, \iota_m \rangle$ of system/entity identifier pairs,
- a timeout specification t_0 of the multi-asset contract expressed in terms of the block number of the atomicity partition,
- the status flag $\mathsf{status} \in \{0, 1\}$ of the contract (initially set to 0), where 1 means that the transaction is completed, and 0 means that it is not yet completed,
- the set $\mathsf{confirmed}$ of already confirmed transactions (initially $\emptyset$).

The transactions of the atomicity partition are:

- reg – registering a new atomic multi-asset swap transaction with contract identifier ι
- con – confirming an existing multi-asset swap transaction with contract identifier ι

The 3-Phase Commit Protocol

Phase 1: Preparation The involved parties prepare transaction orders $P_1, \ldots, P_m$ (in their wallets) that transfer the ownerships of the entities $\iota_1, \ldots, \iota_m$ to special parametrized ownership predicates

$$\varphi_{\mathsf{ato}}(\alpha_0, \iota, t_0, \varphi_1, \varphi'_1; \_;\_,\_,\_)$$
$$,\ldots,$$
$$\varphi_{\mathsf{ato}}(\alpha_0, \iota, t_0, \varphi_m, \varphi'_m; \_;\_,\_,\_).$$

It is not necessary to specify how the parties agree on the terms of the multi-asset swap transaction and which communication channels they use for that. This depends on the domain logic of the involved partitions.

The contract identifier ι is computed as a deterministic pseudo-random function on the transaction orders $P_1, \ldots, P_m$ (without signatures, i.e., ownership proofs) and the ownership predicates

$$\varphi_{\mathsf{ato}}(\alpha_0, \iota, t_0, \varphi_i, \varphi'_i; \_;\_,\_,\_).$$

Definition 9 (Ownership Predicate φ_{ato}). *The predicate $\varphi_{\mathsf{ato}}(\alpha_0, \iota, t_0, \varphi, \varphi'; P; \mathsf{status}, \Pi, s)$ (where the triple $\langle \mathsf{status}, \Pi, s \rangle$ represents the ownership proof), is true if either:*

- *$\varphi'(P, s) = 1$, $\mathsf{status} = 1$, and Π is a certificate in the partition α_0 in a round with number $t < t_0$ of the status of contract ι, or*

– $\varphi(P, s) = 1$, $\mathsf{status} = 0$, *and* Π *is a certificate in the partition* α_0 *in a round with number* $t \geqslant t_0$ *of the status of contract* ι*, or*
– $\varphi(P, s) = 1$, *and* Π *is a certificate in the partition* α_0 *in a round with number* $t \geqslant t_0$ *of the non-existence of contract* ι

The first condition in Def. 9 means that, if all transactions $P_1, \ldots, P_m$ are confirmed in time, the new owners have control over the entities.

The second condition means that, if the status of the contract is still 0 at t_0, the previous owners have control over the entities.

The third condition means that, if the multi-asset swap transaction has not been registered in the atomicity partition at t_0, the previous owners have control over the entities.

Phase 2: Registration One of the parties creates the registration message

$$\langle \alpha_0, \mathsf{reg}, \iota, A, T_0 \rangle,$$

consisting of:

– α_0 – the identifier of the atomicity partition
– reg – the message type
– ι – a contract identifier
– A – a list of attributes that contains the identifiers $\alpha_1, \ldots, \alpha_m, \iota_1, \ldots, \iota_m$ and the timeout t_0 of the multi-asset swap transaction
– T_0 – a timeout specification of the registration message expressed in terms of the block number of the atomicity partition.

The party sends the registration message to the atomicity partition. After it has received the message, the atomicity partition creates a new entity with identifier ι and with the data part

$$D = \langle \alpha_1, \ldots, \alpha_m, \iota_1, \ldots, \iota_m, t_0, \mathsf{status}{=}0, \mathsf{confirmed}{=}\emptyset \rangle$$

.

Phase 3: Confirmation All parties sign their transaction orders $P_1, \ldots, P_m$ (by adding ownership proofs s_i to P_i) and send them to the corresponding partitions $\alpha_1, \ldots, \alpha_m$.

The parties wait until the next blocks in the partitions $\alpha_1, \ldots, \alpha_m$ have been formed and obtain the ledger certificates (of acceptance) $\Pi_1, \ldots, \Pi_m$ for the signed transactions $\langle P_1, s_1 \rangle, \ldots, \langle P_m, s_m \rangle$.

Next, each party k sends the confirmation message

$$\langle \alpha_0, \mathsf{con}, \iota, A_k, T_0 \rangle$$

to the atomicity partition, where

$$A_k = \langle P_k, s_k, \Pi_k \rangle.$$

Upon receiving a confirmation message

$$\big\langle \alpha_0, \mathsf{con}, \iota, \langle P_k, s_k, \Pi_k \rangle, T_0 \big\rangle,$$

the atomicity partition verifies $\langle P_k, s_k, \Pi_k \rangle$, and checks whether the identifiers α_k and ι_k are consistent with the data part

$$D = \langle \alpha_1, \ldots, \alpha_m, \iota_1, \ldots, \iota_m, t_0, \mathsf{status}, \mathsf{confirmed} \rangle$$

of the contract ι. After a successful verification, the triple $\langle P_i, s_i, \Pi_i \rangle$ is added to the set $\mathsf{confirmed}$.

If all transactions have been confirmed, i.e., if

$$\mathsf{confirmed} = \{(P_1, s_1, \Pi_1), \ldots, (P_m, s_m, \Pi_m)\},$$

the atomicity partition checks whether the contract identifier ι was correctly computed on the basis of $P_1, \ldots, P_m$, and finally sets the status flag $\mathsf{status} \leftarrow 1$.

State Tree of the Atomicity Partition The state tree of the atomicity partition is similar to the state tree of Alphabill Money, except that instead of value v there is a state data record

$$D = (\alpha_1, \ldots, \alpha_m, \iota_1, \ldots, \iota_m, t_0, \mathsf{status}, \mathsf{confirmed}),$$

there is neither a summary value V nor ownerships of entities. Hence, the computation of the node hash is given by the formula

$$h \leftarrow H(\iota, H(D, x); h_L; h_R).$$

Implementations with Public Blockchain Technology Basically, the Alphabill platform can be implemented as a private (permissioned) as well as a public (permissionless) blockchain.

Due to the use of certified communication between the partitions and shards of each partition, there are some restrictions to the Alphabill blockchain consensus protocol. There are two categories of consensus protocols used in permissionless blockchains:

- *Persistent consensus* protocols, where a certified n-th block will, once created, be persistently stay in the blockchain. For example, proof-of-stake (PoS), multi-signature-based consensus mechanisms belong to this category.
- *Nakamoto consensus* protocols, where a certified n-th block will not necessarily stay in the blockchain, but instead may be replaced with an alternative block with stronger certificate – although the probability that the n-th block is replaced by a different block after the $(n + k)$-th block has been created becomes negligible when k is large. For example, proof-of-work (PoW) consensus mechanisms belong to this category.

If certified (with ledger certificates) information from the n-th block of one partition (or shard) is used as input to another partition (or shard), then this n-th block must be guaranteed to stay in the first partition (or shard). Otherwise, the certified information might be incorrect and should not be used in other partitions and shards.

Therefore, in case the partitions use independent consensus protocols, only persistent consensus protocols are suitable for the Alphabill platform.

7 Conclusion

We have argued that *uncapped scalability* and *universal tokenization* are the *sine qua non* game changers for blockchain technology to become as impactful as claimed by the plethora of diverse blockchain technology visions stated so far over the last decade, ranging from decentralized (micro-) payment systems over DeFi to now Web3 – to name a few (not yet to speak about the many domain-specific blockchain-based technology visions).

In this paper, we contributed:

- A new form of electronic money scheme, the *bill scheme*, which unlocks, through its decomposability results, unlimited scalability in both permissioned and permissionless scenarios.
- KSI Cash, a central bank digital currency that implements the bill scheme (as outcome of a research cooperation between the Estonian Central Bank and Guardtime) as technological feasibility study of a digital euro. In particular, we have contributed:
  - The design of a bill-based data structure that is optimized for performance. We have provided analytical performance estimations for this data structure.
  - The design of an exchange service, based on a *pro-active* and on a *reactive* exchange strategy.
- Exhaustive performance evaluations of KSI Cash, conducted with the European Central Bank (together with a group of eight central banks from the Eurosystem) encompassing:
  - *A Realistic Backend System Implementation.* Particular effort has been invested to enable realistic usage test. The backend system has been implemented with an effort of 10 person years resulting into approx. 100.000 LOCs.

- *Realistic Usage Tests.* We have conducted a series of test runs under simulated real-world conditions. Most importantly, we conducted a *maximum throughput test run*, showing the system operating with 15.000 transactions per second. (further test runs have been conducted to analyze round calculation times, memory consumption and the behavior of exchange services).
- *Ultra-Scalability Tests.* We have conducted a series of ultra-scalability tests under limited (laboratory) conditions, showing the system managing a load of about 2 million payment orders per second (meaning an equivalent of more than 300.000 transactions per second), demonstrating the linear scalability of KSI Cash (in terms of compute resources)
- *Carbon Footprint Estimation.* We are able to estimate the carbon footprint of KSI Cash as 0.0001g CO2 per transaction, again under the realistic usage scenario (Bitcoin = 100 kg and more)

– We have described the architecture of the Alphabill platform, which enables universal asset tokenization. We have defined the platform's elements, asset presentations, ledger certificates, transaction orders, sharding schemes and transaction system specifications.
– We have introduced Alphabill Money, which is the genuine money partition of the Alphabill platform, as an extended bill scheme.
– We have specified the *dust collection* mechanism of Alphabill Money by elaborating the bill swap scenario and defining the dust collection process.
– *The Alphabill Swap Partition.* Based on the specification of the dedicated Alphabill swap partition, we have defined a 3-phase-commit protocol for atomic heterogeneous (multi-asset) predicate-based swap transactions.

In the design of the Alphabill platform, we have followed the design principles of security-by-design, scalability-by-design, robustness-by-design and viability-by-design. As the consequence of this, with the Alphabill platform, we are able to provide a universal tokenization platform that allows for universal asset tokenization, transfer and exchange as a global medium of exchange.

References

1. Nakamoto, S.: Bitcoin: A Peer-to-Peer Electronic Cash System (2008) [last accessed: 20 April 2022] `https://bitcoin.org/bitcoin.pdf`.
2. Szabo, N.: Smart Contracts: Building Blocks for Digital Markets. Nick Szabo (1996)
3. Szabo, N.: Formalizing and securing relationships on public networks. First Monday **2**(9) (1997)
4. Dixit, A., Deval, V., Dwivedi, V., Norta, A., Draheim, D.: Towards user-centred and legally relevant smart-contract development: A systematic literature review. Journal of Industrial Information Integration **26**(100314) (2022) 1–18
5. Dwivedi, V., Pattanaik, V., Deval, V., Dixit, A., Norta, A., Draheim., D.: Legally enforceable smart-contract languages: A systematic literature review. ACM Computing Surveys **54**(5) (2022) 1–34
6. Grassi, L., Lanfranchi, D., Faes, A., Renga, F.M.: Do we still need financial intermediation? the case of decentralized finance – DeFi. Qualitative Research in Accounting & Management (February 2022) 1–22
7. Visa USA: Visa acceptance for retailers [last accessed: 11 March 2022] `https://usa.visa.com/run-your-business/small-business-tools/retail.html`.
8. Meyer, F., Kuhlmann, L., Müller, S.: Card schemes – Europe could irreversibly lose its position in the payment market. CORE SE (2019) [last accessed: 11 March 2022] `https://core.se/techmonitor/card-schemes`.
9. Yang, Z., Yang, R., Yu, F.R., Li, M., Zhang, Y., Teng, Y.: Sharded blockchain for collaborative computing in the Internet of Things: Combined of dynamic clustering and deep reinforcement learning approach. IEEE Internet of Things Journal (2022) 1–17
10. Hameed, S., Shah, S.A., Saeed, Q.S., Siddiqui, S., Ali, I., Vedeshin, A., Draheim, D.: A scalable key and trust management solution for IoT sensors using SDN and blockchain technology. IEEE Sensors **21**(6) (2021) 8716–8733
11. Vedeshin, A., Dogru, J.M., Liiv, I., Ben Yahia, S., Draheim, D.: A secure data infrastructure for personal manufacturing based on a novel key-less, byte-less encryption method. IEEE Access **8** (2020) 40039–40056
12. Vedeshin, A., Dogru, J.M.U., Liiv, I., Draheim, D., Ben Yahia, S.: A digital ecosystem for personal manufacturing: An architecture for a cloud-based distributed manufacturing operating system. In: Proceedings of MEDES'2019 – the 11th International Conference on Management of Digital EcoSystems, ACM (2019) 224–228

13. Sagirlar, G., Carminati, B., Ferrari, E., Sheehan, J.D., Ragnoli, E.: Hybrid-IoT: Hybrid blockchain architecture for Internet of Things – PoW sub-blockchains. In: Proceedings of iThings'2018: the 11th IEEE International Conference on Internet of Things, IEEE (2018) 1007–1016
14. Buldas, A., Draheim, D., Nagumo, T., Vedeshin, A.: Blockchain technology: Intrinsic technological and socio-economic barriers. In: Proceedings of FDSE'2020 – the 7th International Conference on Future Data and Security Engineering. Volume 12466 of Lecture Notes in Computer Science., Springer (2020) 3–27
15. Williamson, O.E.: Transaction cost economics: How it works; where it is headed. De Economist **146** (1998) 23–58
16. Rikken, O., Janssen, M., Kwee, Z.: Governance challenges of blockchain and decentralized autonomous organizations. Information Polity **24**(4) (2019) 397–417
17. Eesti Pank, Bank of Greece, Deutsche Bundesbank, Central Bank of Ireland, Banco de España, Latvijas Banka, European Central Bank, Banca d'Italia, De Nederlandsche Bank: Work Stream 3: A New Solution – Blockchain & eID. July 2021. [last accessed: 28 March 2022] `https://www.ecb.europa.eu/paym/digital_euro/investigation/profuse/shared/files/deexp/ecb.deexp211011_3.en.pdf` [last accessed: 28 March 2022] `https://haldus.eestipank.ee/sites/default/files/2021-07/Workstream3-ANewSolution-BlockchainandeID_1.pdf`.
18. Olt, R., Meidla, T., Ilves, L., Steiner, J.: Summary report: Results of the Eesti Pank - Guardtime CBDC Research. Eesti Pank, Guardtime (December 2021) [last accessed: 11 March 2022] `https://haldus.eestipank.ee/sites/default/files/2021-12/EP-Guardtime_CBDC_Research_2021_eng.pdf`.
19. Buldas, A., Saarepera, M., Steiner, J., Ilves, L., Olt, R., Meidla, T.: Formal Model of Money Schemes and their Implications for Central Bank Digital Currency. Eesti Pank, Guardtime (2021) [last accessed: 11 March 2022] `https://haldus.eestipank.ee/sites/default/files/2021-12/EP-A_Formal_Model_of_Money_2021_eng.pdf`.
20. Buldas, A., Saarepera, M., Steiner, J., Draheim, D.: A Unifying Theory of Electronic Money and Payment Systems. TechRxiv (2021) `https://doi.org/10.36227/techrxiv.14994558.v1`.
21. International Organization for Standardization: ISO 22739:2020(E) – Blockchain and distributed ledger technologies – Vocabulary. ISO (2020)
22. Yaga, D., Mell, P., Roby, N., Scarfone, K.: Blockchain technology overview. Technical Report NISTIR 8202, National Institute of Standards and Technology, Gaithersburg (October 2018)
23. Herrera-Joancomartí, J., Pérez-Solà, C.: Privacy in Bitcoin transactions: New challenges from blockchain scalability solutions. In Torra, V., Narukawa, Y., Navarro-Arribas, G., Yañez, C., eds.: Proceedings of MDAI'2016 – the 13th International Conference on Modeling Decisions for Artificial Intelligence. Volume 9880 of Lecture Notes in Artificial Intelligence., Springer (2016) 26–44
24. Decker, C., Wattenhofer, R.: A fast and scalable payment network with Bitcoin duplex micropayment channels. In Pelc, A., Schwarzmann, A.A., eds.: Proceedings of SSS'2015 – the 17th International Symposium on Stabilization, Safety, and Security of Distributed Systems. Volume 9212 of Lecture Notes in Computer Science., Springer (2015) 3–18
25. Poon, J., Dryja, T.: The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments, Draft Version 0.5.9.2 (January 14, 2016) [last accessed: 18 April 2022] `https://lightning.network/lightning-network-paper.pdf`.
26. Xie, J., Yu, F.R., Huang, T., Xie, R., Liu, J., Liu, Y.: A survey on the scalability of blockchain systems. IEEE Network **33**(5) (2019) 166–173
27. Luu, L., Narayanan, V., Zheng, C., Baweja, K., Gilbert, S., Saxena, P.: A secure sharding protocol for open blockchains. In Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S., eds.: Proceedings of CCS'16 – the 23rd ACM SIGSAC Conference on Computer and Communications Security, New York, NY, USA, Association for Computing Machinery (2016) 17–30
28. Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., Ford, B.: OmniLedger: A secure, scale-out, decentralized ledger via sharding. Cryptology ePrint Archive **Report 2017/406** (11 May 2017) 1–16
29. Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., Ford, B.: OmniLedger: A secure, scale-out, decentralized ledger via sharding. In: Proceedings of S&P'18 – the 39th IEEE Symposium on Security and Privacy, IEEE (2018) 583–598
30. Yu, G., Wang, X., Yu, K., Ni, W., Zhang, J.A., Liu, R.P.: Survey: Sharding in blockchains. IEEE Access **8** (2020) 14155–14181
31. Hafid, A., Hafid, A.S., Samih, M.: Scaling blockchains: A comprehensive survey. IEEE Access **8** (2020) 125244–125262
32. Al-Bassam, M., Sonnino, A., Bano, S., Hrycyszyn, D., Danezis, G.: Chainspace: A sharded smart contracts platform. CoRR **abs/1708.03778** (2017) 1–16
33. Al-Bassam, M., Sonnino, A., Bano, S., Hrycyszyn, D., Danezis, G.: Chainspace: A sharded smart contracts platform. In: Proceedings of NDSS'18 – the 25th Annual Network and Distributed System Security Symposium, The Internet Society (2018)

34. Zamani, M., Movahedi, M., Raykova, M.: RapidChain: Scaling blockchain via full sharding. In Lie, D., Mannan, M., Backes, M., Wang, X., eds.: Proceedings of CCS'18 – the 25th ACM SIGSAC Conference on Computer and Communications Security, New York, NY, USA, Association for Computing Machinery (2018)
35. Wang, J., Wang, H.: Monoxide: Scale out blockchains with asynchronous consensus zones. In: Proceedings of NSDI'19 – 16th USENIX Symposium on Networked Systems Design and Implementation, USENIX Association (2019) 95–112
36. Drake, J.: Ethereum sharding (17 May 2018) [last accessed: 26 March 2022] `https://youtu.be/J4rylD6w2S4`.
37. Dang, H., Dinh, T.T.A., Loghin, D., Chang, E.C., Lin, Q., Ooi, B.C.: Towards scaling Blockchain systems via sharding. In: Proceedings of SIGMOD'19: the 2019 International Conference on Management of Data, New York, NY, USA, Association for Computing Machinery (2019) 123–140
38. Chen, H., Wang, Y.: SSChain: A full sharding protocol for public blockchain without data migration overhead. Pervasive and Mobile Computing **59**(101055) (October 2019) 1–15
39. Manuskin, A., Mirkin, M., Eyal, I.: Ostraka: Secure blockchain scaling by node sharding. In: Proceedings of Euro S&PW'2020 – the 2020 IEEE European Symposium on Security and Privacy Workshops, IEEE (2020) 397–406
40. Kokoris-Kogias, E., Jovanovic, P., Gailly, N., Khoffi, I., Gasser, L., Ford, B.: Enhancing Bitcoin security and performance with strong consistency via collective signing. In: Proceedings of USENIX Security'16 – the 25th USENIX Security Symposium, USENIX Association (2016) 279–296
41. Pass, R., Shi, E.: Hybrid consensus: Efficient consensus in the permissionless model. Cryptology ePrint Archive, Report 2016/917 (2016) `https://ia.cr/2016/917`.
42. Pass, R., Shi, E.: Hybrid consensus: Efficient consensus in the permissionless model. In Richa, A.W., ed.: Proceedings of DISC'2017 – the 31st International Symposium on Distributed Computing. Volume 91 of Leibniz International Proceedings in Informatics (LIPIcs)., Dagstuhl, Germany, Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2017) 39:1–39:16
43. Lamport, L., Shostak, R., Peas, M.: The Byzantine generals problem. ACM Transactions on Programming Languages and Systems **4**(3) (July 1982) 382–401
44. Kiayias, A., Russell, A., David, B., Oliynykov, R.: Ouroboros: A provably secure proof-of-stake blockchain protocol. In Katz, J., Shacham, H., eds.: Proceedings of CRYPTO'17 – 17th Annual International Cryptology Conference, Part I. Volume 10401 of Lecture Notes in Computer Science., Springer (2017) 357–388
45. Gilad, Y., Hemo, R., Micali, S., Vlachos, G., Zeldovich, N.: Algorand: Scaling Byzantine agreements for cryptocurrencies. In: Proceedings of SOSP'17 – the 26th Symposium on Operating Systems Principles, ACM (October 2017) 51–68
46. Syta, E., Jovanovic, P., Kogias, E.K., Gailly, N., Gasser, L., Khoffi, I., Fischer, M.J., Ford, B.: Scalable bias-resistant distributed randomness. In: Proceedings of SP'2017 – the 38th IEEE Symposium on Security and Privacy, IEEE (2017) 444–460
47. Castro, M., Liskov, B.: Practical Byzantine fault-tolerance and proactive recovery. ACM Transactions Computer Systems **20**(4) (2002) 398–461
48. Sousa, J., Bessani, A.N.: From byzantine consensus to BFT state machine replication: A latency-optimal transformation. In: Proceedings of EDCC'12 – the 2012 9th European Dependable Computing Conference, IEEE Computer Society (2012) 37–48
49. Buterin, V.: A next generation smart contract and decentralized application platform – Ethereum white paper (2015)
50. Okanami, N., Nakamura, R., Nishide, T.: Load balancing with in-protocol/wallet-level account assignment in sharded blockchains. IEICE Transactions on Information and Systems **E105D**(2) (2022) 205–214
51. Ren, L., Nayak, K., Abraham, I., Devadas, S.: Practical synchronous byzantine consensus. CoRR **abs/1704.02397** (2017)
52. Alon, N., Kaplan, H., Krivelevich, M., Malkhi, D., Stern, J.: Addendum to "Scalable secure storage when half the system is faulty". Information and Computation **205**(7) (July 2007) 1114–1116
53. Maymounkov, P., Mazières, D.: Kademlia: A peer-to-peer information system based on the XOR metric. In Druschel, P., Kaashoek, F., Rowstron, A., eds.: Proceedings of IPTPS'2002 – the 1st International Workshop on Peer-to-Peer Systems. Volume 2429 of Lecture Notes in Computer Science., Berlin, Heidelberg, Springer (2002) 53–65
54. Hafid, A., Hafid, A.S., Samih, M.: A novel methodology-based joint hypergeometric distribution to analyze the security of sharded blockchains. IEEE Access **8** (2020) 179389–179399
55. Du, M., Chen, Q., Ma, X.: MBFT: A new consensus algorithm for consortium blockchain. IEEE Access **8** (2020) 87665–87675
56. Sohrabi, N., Tari, Z.: ZyConChain: A scalable blockchain for general applications. IEEE Access **8** (2020) 158893–158910

57. Kotla, R., Alvisi, L., Dahlin, M., Clement, A., Wong, E.: Zyzzyva: Speculative byzantine fault tolerance. ACM Transactions on Computer Systems **27**(4) (2010)
58. Jia, D., Xin, J., Wang, Z., Wang, G.: Optimized data storage method for sharding-based blockchain. IEEE Access **9** (2021) 67890–67900
59. Huang, H., Yue, Z., Peng, X., He, L., Chen, W., Dai, H.N., Zheng, Z., Guo, S.: Elastic resource allocation against imbalanced transaction assignments in sharding-based permissioned blockchains. IEEE Transactions on Parallel and Distributed Systems **33**(10) (2022) 2372–2385
60. Neely, M.: Stochastic Network Optimization with Application to Communication and Queueing Systems. Morgan & Claypool (2010)
61. Georgiadis, L., Neely, M.J., Tassiulas, L.: Resource allocation and cross-layer control in wireless networks. Foundations and Trends in Networking **1**(1) (2006) 1–144
62. Danezis, G., Meiklejohn, S.: Centrally banked cryptocurrencies. In: Proceedings of NDSS'16 – the 23rd Annual Symposium on Network and Distributed System Security, The Internet Society (2016) 1–14
63. Federal Reserve Bank of Boston and Massachusetts Institute of Technology Digital Currency Initiative: Project Hamilton Phase 1 – A High Performance Payment Processing System Designed for Central Bank Digital Currencies. Federal Reserve Bank of Boston (3 February 2022) [last accessed: 26 March 2022] `https://www.bostonfed.org/-/media/Documents/Project-Hamilton/Project-Hamilton-Phase-1-Whitepaper.pdf`.
64. Zhong, L., Wu, Q., Xie, J., Guan, Z., Qin, B.: A secure large-scale instant payment system based on blockchain. Computers & Security **84** (2019) 349–364
65. Wood, G.: Polkadot: Vision for a heterogneous multi-chain framework, Draft 1 (2016) [last accessed: 3 Feb 2022] `https://polkadot.network/PolkaDotPaper.pdf`.
66. Beer, S.: The viable system model: Its provenance, development, methodology and pathology. Journal of the Operational Research Society **35**(1) (1984) 7–25
67. The Economist: The real-time revolution. The Economist October 23rd-29th, pages 22–24, 2021
68. The Economist: Instant economocis. The Economist October 23rd-29th, page 13, 2021
69. Beer, S.: Brain Of The Firm. Allen Lane, the Penguin Press (1972)
70. Beer, S.: The Heart of the Enterprise. John Wiley & Sons (1979)
71. Holland, J.H.: Studying complex adaptive systems. Journal of Systems Science and Complexity **19**(1) (2006) 1–8
72. Miller, J.H., Page, S.E.: Complex Adaptive Systems: An Introduction to Computational Models of Social Life. Princeton University Press (2007)
73. Latour, B.: Reassembling the Social: An Introduction to Actor–Network Theory. Oxford University Press (2005)
74. McBride, K., Draheim, D.: On complex adaptive systems and electronic government – a proposed theoretical approach for electronic government studies. The Electronic Journal of e-Government **18**(1) (2020) 43–53
75. Eesti Pank: Eesti Pank ran an experiment to investigate the technological possibilities of a central bank digital currency based on blockchain, Eesti Pank, 13 December 2021 `https://tinyurl.com/2djdvmn9`.
76. Antonopoulos, A.M.: Mastering Bitcoin: Programming the Open Blockchain. O'Reilly (2017)
77. Merkle, R.C.: Protocols for public key cryptosystems. In: Proceedings of S&P 1980 – the 1st IEEE Symposium on Security and Privacy, IEEE Computer Society (1980) 122–134
78. Merkle, R.C.: Method of Providing Digital Signatures. US Patent No.: US4309569A (1982)
79. Merkle, R.C.: A digital signature based on a conventional encryption function. In Pomerance, C., ed.: Proceedings of CRYPTO'1987 – the 7th Conference on the Theory and Applications of Cryptographic Techniques. Volume 293 of Lecture Notes in Computer Science., Berlin, Heidelberg, Springer (1987) 369–378
80. Accredited Standards Committee X9: American National Standard X9.62-2005, Public Key Cryptography for the Financial Services Industry, The Elliptic Curve Digital Signature Algorithm (ECDSA) (November 2005)
81. Rasmussen, N.: Calculating total cooling requirements for data centers. Technical Report APC White Paper 25, revision 3, Schneider Electric's Data Center Science Center (2017)
82. Foteinis, S.: Bitcoin's alarming carbon footprint. Nature **554** (February 2018) 169
83. Sandner, P., Lichti, C., Heidt, C., Richter, R., Schaub, B.: The Carbon Emissions of Bitcoin From an Investor Perspective. Frankfurt School Blockchain Center (2021)
84. Trespalacios, J.P., Dijk, J.: The Carbon Footprint of Bitcoin. De Nederlandsche Bank (2021) [last accessed: 29 March 2022] `https://www.dnb.nl/media/1ftd2xjl/the-carbon-footprint-of-bitcoin.pdf`.
85. Stoll, C., Klaaßen, L., Gallersdörfer, U.: The carbon footprint of Bitcoin. Technical Report CEEPR WP 2018-018, MIT Center for Energy and Environmental Policy Research, Cambridge, MA (December 2018)
86. Stoll, C., Klaaßen, L., Gallersdörfer, U.: The carbon footprint of Bitcoin. Joule **3**(7) (2019) 1647–1661
87. Badea, L., Mungiu-Pupazan, M.C.: The economic and environmental impact of Bitcoin. IEEE Access **9** (2021) 48091–48104

88. Buldas, A., Laud, P., Lipmaa, H.: Eliminating counterevidence with applications to accountable certificate management. Journal of Computer Security **10**(3) (September 2002) 273–296
89. Buldas, A., Laur, S.: Knowledge-binding commitments with applications in time-stamping. In: Proceedings of PKC'2007 – the 10th International Conference on Practice and Theory in Public-Key Cryptography. Volume 4450 of Lecture Notes in Computer Science. (2007) 150–165
90. Emmelhainz, M.A.: EDI: A Total Management Guide. Van Nostrand Reinhold (1993)
91. Adelson-Velsky, G., Landis, E.: An algorithm for the organization of information (in Russian). Doklady Akademii Nauk (Proceedings of the Russian Academy of Sciences) **146**(2) (1962) 263–266

Ahto Buldas is professor of cryptography at Tallinn University of Technology. Ahto studied computer science at Tallinn University of Technology (1985-1991) and holds an MSc on simulation techniques for Boolean circuits (1992) and a PhD on computational algebraic graph theory (1999). Ahto's research interests are related to applied cryptography. His time-stamping related research started in 1997 and he has published papers in the conferences Crypto, Asiacrypt and PKC. Ahto participated in the development of the Estonian Digital Signature Act and the Estonina eID card (1996-2002). His current research interests also include risk analysis methods, including attack-tree semantics and game-theoretical approaches to risk analysis. Ahto Buldas is a co-founder of Guardtime and also of Cybernetica AS.

Dirk Draheim received the PhD from Freie Universität Berlin and the habilitation from Universität Mannheim, Germany. Currently, he is full professor of information society technology at Tallinn University of Technology and head of the Information Systems Group, Tallinn University of Technology, Estonia. The Information Systems Group conducts research in large and ultra-large-scale IT systems. He is also an initiator and leader of numerous digital transformation initiatives. Dirk is author of the Springer books "Business Process Technology", "Semantics of the Probabilistic Typed Lambda Calculus" and "Generalized Jeffrey Conditionalization", and co-author of the Springer book "Form-Oriented Analysis".

Mike Gault received the PhD from the University of Wales in 1994. Mike is the CEO and founder of Guardtime, a company he has led for the last 10 years. He started his career on a post-doctoral fellowship conducting research at the Tokyo Institute of Technology on the numerical analysis of quantum devices. He then spent 10 years as a quant and derivatives trader at Credit Suisse Financial Products and Barclays Capital before starting Guardtime.

Risto Laanoja is pursuing his PhD degree at Tallinn University of Technology, working on provable security of hashing based constructs and the development of "strong pre-image awareness" and "bounded pre-image awareness" concepts. He is a member of the R&D team at Guardtime.

Takehiko Nagumo is Senior Managing Executive Officer of Mitsubishi UFJ Research and Consulting. Also, Takehiko is adjunct professor in strategic management at Kyoto University Graduate School of Management. Takehiko holds an MSc in Development Finance from University of London and an MBA in Strategic Management from Georgetown University. Numerous times, Takehiko's research on Balanced Scorecard has been published by Harvard Business School. Takehiko is one of the thought leaders in Japan in the field of digital economy and society: he is a member of the Japanese Government's Regulatory Reform Promotion Committee, research fellow at the National Institute of Advanced Industrial Science and Technology, co-founder and Executive Director of the Smart City Institute Japan and Fellow of the World Economic Forum, Centre for Fourth Industrial Revolution Japan.

Märt Saarepera has extraordinary experience in developing blockchain systems. Also, he is an experienced business developer, having shown his leadership skills by growing small startups to strong enterprises. Märt was a founder of Guardtime in 2006, which has now grown to one of the leading industrial blockchain developing companies providing reliable data integrity assurance solutions for industry and the public sector.

Märt holds a doctoral degree in real-time systems verification from the Tokyo Institute of Technology and has been co-author of several research papers on cryptography, integrity protection and privacy.

Syed Attique Shah received the Ph.D. degree from the Institute of Informatics, Istanbul Technical University, Istanbul, Turkey. During his Ph.D., he studied as a Visiting Scholar at the University of Tokyo, Japan, the National Chiao Tung University, Taiwan, and the Tallinn University of Technology, Estonia, where he completed the major content of his thesis. He has worked as an Associate Professor and the Chairperson at the Department of Computer Science, BUITEMS, Quetta, Pakistan. He was also engaged as a Lecturer at the Data Systems Group, Institute of Computer Science, University of Tartu, Estonia. Currently, he is working as a Lecturer in Smart Computer Systems, at the School of Computing and Digital Technology, Birmingham City University, United Kingdom. His research interests include big data analytics, the Internet of Things, network security, and information management.

Joosep Simm is software engineer at Guardtime. He studied computer science at Tallinn University of Technology (2002-2010). Joosep has been active in software development from 2005.

Jamie Steiner has 25 years of experience in Intelligence, Finance, and Technology. While in the US Air Force, he collected and analyzed signals and imagery intelligence, and crafted the Air Force's premier course on precise geographical positioning for targeting. His experience in finance includes making markets in inflation linked government bonds for JPMorgan Chase, overhauling the fixed income trading systems used by the JPMorgan Private Bank, and managing regulatory and compliance issues. As the Head of Electronic Trading at Phoenix Partners LLC, he designed, tested, and managed one of the top electronic liquidity venues for credit default swaps. Jamie received a bachelor's degree in Aeronautical Engineering from the United States Air Force Academy, and a master's degree in Business Administration from the NYU Stern School of Business.

Tanel Tammet is a full professor of applied artificial intelligence at Tallinn University of Technology. His main research interests include crowd-sourced knowledge bases, automated reasoning and commonsense reasoning. He has also worked in the area of cybersecurity and has been involved in numerous large commercial and public sector IT projects. Tanel was one of the initiators of the Estonian X-ROAD interoperability framework and has helped to develop several other core IT infrastructure systems in Estonia.

Ahto Truu is a software architect and a member of the R&D team at Guardtime. He holds an MSc in computer science from Tartu University and a PhD from Tallinn University of Technology, where his thesis was on hash-based server-assisted digital signature solutions.