FinTech

MDPI

Article

Pravuil: Global Consensus for a United World

David Cerezo Sánchez

Calctopia, T23 AT2P Cork, Ireland; david@calctopia.com

Abstract: The latest developments in blockchain technology have conceptualised very efficient consensus protocols that have not yet been able to overcome older technologies. This paper presents Pravuil, a robust, secure, and scalable consensus protocol for a permissionless blockchain suitable for deployment in an adversarial environment such as the Internet. Using zero-knowledge authentication techniques, Pravuil circumvents previous shortcomings of other blockchains: Bitcoin's limited adoption problem (as transaction demand grows, payment confirmation times grow much less than that of other PoW blockchains); higher transaction security at a lower cost; more decentralisation than other permissionless blockchains; impossibility of full decentralisation; the blockchain scalability trilemma (decentralisation, scalability, and security can be achieved simultaneously); and Sybil resistance for free implementation of the social optimum. Pravuil goes beyond the economic limits of Bitcoin and other PoW/PoS blockchains, leading to a more valuable and stable cryptocurrency.

Keywords: consensus; permissionless; permissioned; scalability; zero-knowledge; mutual attestation

check for updates

Citation: Cerezo Sánchez, D. Pravuil: Global Consensus for a United World. *FinTech* **2022**, *1*, 325–344. https://doi.org/10.3390/fintech1040025

Academic Editor: David Roubaud

Received: 3 September 2022
Accepted: 27 October 2022
Published: 31 October 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2022 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

A third generation of blockchains has been developed featuring the latest advances in cryptography and sharding to reach maximum performance and security in Internet settings: they usually make use of advances in BFT-like consensus protocols [1,2] and collective signatures [3] to obtain 1000s of transactions per second. Despite these impressive achievements in modern distributed computing, we are motivated to find solutions to the lack of lawful, green blockchains that are also at the edge of efficiency according to the latest research in economics (see Section 5 for an in-depth discussion).

In this work, we introduce Pravuil (i.e., in the Book of the Secrets of Enoch, an archangel "swifter in wisdom than the other archangels", scribe and recordkeeper -analogous to the use of consensus protocol to keep records of agreements-), a robust, secure, and scalable consensus protocol for real-world deployments on open, permissionless environments that, unlike other proposals, remains robust to high adversarial power and adaptation while considering rational participants and providing strong consistency (e.g., no forks, forward-security, and instant transactions). Our protocol is also the first to integrate real-world identity on Layer 1, as required by current financial regulations, obtaining Sybil resistance for free: a very useful property considering the electrical waste produced by Bitcoin, its Achilles' heel, which this blockchain circumvents for the first time by obviating to pay the Price of Crypto-Anarchy [4].

To achieve the desired goals, we introduce a new consensus protocol in which we prioritise robustness against attackers and censorship. We then incorporate zero-knowledge proof-of-identity [4] while maintaining an open, permissionless node membership mechanism enabling high levels of decentralisation.

Contributions

In summary, we make the following contributions:

- We propose a consensus protocol that remains robust, secure, and scalable among rational participants in an Internet setting;

- We prove liveness, safety, and censorship-resistance of our new consensus protocol;
- We discuss the underlying rationale of our design and prove all the advantages that it provides over previous blockchain designs.

2. Related Literature

Previous blockchain designs [1–3,5] deal with different trade-offs of the scalability trilemma (security vs. scalability vs. decentralisation), and they are not usually concerned with the economic consequences of their design (e.g., the Price of Crypto-Anarchy [4]) or the legal consequences of the lack of real-world identity as required by recent legislation (the FATF Travel Rule). In this paper, we further extend beyond the usual criteria of consensus literature by incorporating economic and legal considerations as primary motivations to design blockchain protocols; for previous literature analysing multiple blockchains and cryptocurrencies from a general perspective, see [6]; for previous literature comparing the economic consequences of different architectures of blockchains, see [7,8]; for previous literature about the economic consequences of Bitcoin's technical design, see [9–13]; for equilibrium models of cryptocurrencies according to their technical features, see [14–16].

Previous designs of ByzCoin/OmniLedger/MOTOR [17–19] proposed proof-of-work (PoW) as a Sybil-resistance mechanism; although their consensus protocols are more advanced and performant than Bitcoin's (i.e., better scalability in Table 1), they still pay for the Price of Crypto-Anarchy [4] (i.e., not free Sybil resistance in Table 1 and Section 5.6). Further, although other blockchains (e.g., [20]) provide methods to anonymise real-world identities (i.e., real-world identities in Table 1), in order to create lawful blockchains (i.e., lawfulness in Table 1), they fail to incorporate these privacy techniques into their consensus protocol, as they keep on using proof-of-stake as a Sybil-resistance mechanism; thus, they still pay the Price of Crypto-Anarchy [4] (i.e., not free Sybil resistance in Table 1 and Section 5.6), suffer from Bitcoin's limited adoption problem [9] (i.e., unlimited adoption in Table 1 and Section 5.1), and exist within its same economic limits [10] (i.e., no economic limitation in Table 1 and Section 5.7).

Table 1. Features of different consensus protocols.

Features	Bitcoin	ByzCoin/MOTOR	Pravuil
Secure	✓	✓	✓
Decentralised	✓	✓	✓
Scalability	☒	✓	✓
Real-world identity	☒	☒	✓
Free Sybil resistance	☒	☒	✓
Lawfulness	☒	☒	✓
Unlimited adoption	☒	☒	✓
No economic limitation	☒	☒	✓

In order to better understand the lineage of consensus protocols leading up to Pravuil, we briefly recap their history: Bitcoin introduced a novel consensus protocol that was open and permissionless over classical variants by using Nakamoto consensus and proof-of-work (PoW), but its key problems were its probabilistic consistency guarantees and the use of costly PoW as a Sybil-resistance mechanism. ByzCoin improved over Bitcoin by introducing strong consistency (e.g., blocks and transactions are almost instantly valid) and better scalability up to hundreds of nodes using collective signatures; then, OmniLedger introduced sharding over ByzCoin to scale up to many thousands of nodes and parallelise transaction processing; finally, MOTOR improved the robustness of ByzCoin/OmniLedger, making it suitable for open, adversarial networks such as the Internet. However, none of the previous consensus protocols tackled the problem of how to use real-world identities on a blockchain while remaining permissionless: Pravuil introduces anonymised versions of real-world identities, thus obviating the high costs of the Price of Crypto-Anarchy that is endemic on PoW/PoS mechanisms for Sybil resistance. Additional benefits include:

improving lawfulness in accordance with the latest regulations; better economics that removes barriers such as Bitcoin's limited adoption problem and its economic limits; and bootstrapping a payment network with billions of addressable users from its very inception.

3. Model and Methods

In this work, we make use of the following definitions, security model, security assumptions, and consensus protocols.

Definition 1 (Strongly consistent broadcast [21]). *A protocol for strongly consistent broadcast satisfies the following conditions except with negligible probability:*

- *Termination: If a correct party strongly consistently broadcasts m with tag ID, then all correct parties eventually strongly consistently deliver m with tag ID.*
- *Agreement: If two correct parties P_i and P_j strongly consistently deliver m and m', respectively, with tag ID, then $m = m'$.*
- *Integrity: Every correct party strongly consistently delivers at most one payload m with tag ID. Moreover, if the sender P_s is correct, then m was previously strongly consistently broadcast by P_s with tag ID.*
- *Transferability: After a correct party has strongly consistently delivered m with tag ID, it can generate a string M_{ID} such that any correct party that has not strongly consistently delivered message with tag ID is able to strongly consistently deliver some message immediately upon processing M_{ID}.*
- *Strong unforgeability: For any ID, it is computationally infeasible to generate a value M that is accepted as valid by the validation algorithm for completing ID unless $n - 2t$ correct parties have initialised instance ID and actively participated in the protocol.*
- *Improved communication complexity: This reduces the optimistic case to $O(n)$, unlike the Hyperledger Sawtooth PBFT of $O(n^2)$.*

Definition 2 (Partial synchronous model [22,23]). *In a partially synchronous network, there is a known bound Δ and an unknown global stabilisation time (GST), such that after GST, all transmissions between honest nodes arrive within time Δ.*

Definition 3 (n = 3f + 1 [24]). *The proportion of malicious nodes that an adversary controls accounts for no more than 1/3 of the whole shard. The rest of the nodes are rational, that is, maximisers of their transaction rewards.*

Definition 4 (Round-adaptive adversary [25]). *We assume a mildly adaptive, computationally bounded adversary that chooses which nodes to corrupt at the end of every consensus round and has control over them at the end of the next round.*

Definition 5 (Strong consistency [17]). *The generation of each block is deterministic and instant, with the following features:*

- *There is no fork in a blockchain. By running a distributed consensus algorithm, state machine replication is achieved.*
- *Transactions are confirmed almost instantly. Whenever a transaction is written into a block, the transaction is regarded as valid.*
- *Transactions are tamper-proof (forward security). Whenever a transaction is written to a blockchain, the transaction and block cannot be tampered with and the block will remain on the chain at all times.*

Definition 6 (BLS [26] and BDN [27] signatures). *Boneh–Lynn–Sacham (BLS) and Boneh–Drijvers–Neven (BDN) signatures are assumed secure.*

Definition 7 (Global PKI [28]). *Our blockchain design assumes a global public key infrastructure (PKI), not directly for consensus purposes, but as a node-admission and Sybil-resistance mechanism [4].*

Definition 8 (Permissionless network [29]). *In a permissionless network:*

- *Anyone can join a node without requiring permission from any party;*
- *Any node can join or leave at any time;*
- *The number of participating nodes varies at any time and is unpredictable.*

3.1. Prior Consensus Protocols

Pravuil builds over ByzCoin [17], OmniLedger [18], and MOTOR [19]. ByzCoin [17] envisions a Bitcoin [30] protocol that uses strongly consistent consensus, scaling with multi-cast trees, and aggregate Schnorr signatures. OmniLedger [18] adds sharding over ByzCoin [17], and MOTOR [19] strengthens the robustness of ByzCoin [17] for an open, adversarial network such as the Internet.

In the next section, we extend these consensus protocols to address issues that prevent their deployment in an adversarial environment such as the Internet.

3.2. Research Design

Our research design combines different methods:

1. A broad literature review about the economic consequences of the technical features used by blockchains and cryptocurrencies leads us to formulate an optimal architecture; detailed teachings distilled from said literature review are detailed in Section 5;
2. The optimal architecture is described and implemented in Section 4;
3. A controlled experiment simulating a real-world deployment under different configurations is described in Section 6.

4. Resulting Design

Pravuil improves previous work by using another source of randomness, drand [31], and by incorporating zero-knowledge proof-of-identity [4] as a Sybil-resistance mechanism into the first layer of the consensus protocol.

4.1. Goals

To sum up, Pravuil has the following goals:

- Robustness: the consensus round can only be disrupted by controlling the leader node.
- Scalability: the protocol performs well among hundreds of nodes ($n = 600$).
- Fairness: the malicious leader can only be elected with a probability equal to the percentage of malicious nodes in the system (i.e., the adversary cannot always control the leader).

We detail the extensions over a previous Byzantine fault tolerance (BFT) protocol such as ByzCoin/MOTOR in order to obtain an improved blockchain-consensus algorithm.

4.2. Rotating Leader

View-change protocols assume a predetermined schedule of leaders, making them susceptible to adversaries that compromise the next f leaders.

To prevent this attack, our blockchain uses drand [31]: an efficient randomness beacon daemon that utilises bilinear pairing-based cryptography, t-of-n distributed key generation, and threshold BLS [26] signatures to generate publicly verifiable, unbiasable, unpredictable, highly available, distributed randomness at fixed time intervals. As described in its online specification [32], drand uses the BLS12-381 curve, the Feldman [33] verifiable secret sharing protocol, and the joint Feldman protocol [34] for distributed key generation (DKG); using threshold BLS signatures as a source of randomness has been proven to be secure [35] according to its security model [36].

Remark 1. *In this work, we inherit all the previous security theorems from ByzCoin [17], OmniLedger [18], and MOTOR [19]. Note that thanks to BLS collective signatures, signatures get compacted close to $O(1)$, per-round complexity is reduced to $O(logN)$, and signature verification complexity is reduced to $O(1)$, unlike Hyperledger Sawtooth PBFT complexities of $O(N)$, $O(N^2)$, and $O(N)$.*

Theorem 1 (Robustness/Liveness)**.** *The adversary cannot predict nor bias the leader election.*

Proof. The unpredictability property follows from the unforgeability of the BLS [26] signing algorithm, and the unbiasability property follows from the deterministic nature of the BLS [26] signing algorithm. The leader of view v is determined by the outcomes of drand's public service (see NewViewChangeRequest in Algorithm 1), and all the nodes can publicly verify its election when needed (see ViewChangeVerifyRandom in Algorithm 1)drandPseudoCodedrandPseudoCode. Thus, the adversary cannot predict nor bias the leader election, preventing the adversary from breaking liveness. □

Theorem 2 (Safety/Censorship-resistance)**.** *A round-adaptive adversary cannot always control the consensus decision.*

Proof. As the leader election is unpredictable (Theorem 1), the adversary can only hope that one of its randomly compromised nodes gets chosen. Given that

$$\frac{1}{3^d}$$

is the probability that the adversary controls d consecutive leaders, the adversary cannot control the leader forever, since

$$\lim_{d \to \infty} \frac{1}{3^d} = 0.$$

Thus, the adversary cannot always control the consensus decision. □

4.3. Zero-Knowledge Proof-of-Identity

As explained in a paper about the economic limits of Bitcoin [10], Bitcoin is prohibitively expensive to run because the recurring "flow" payments to miners for running the blockchain (particularly, the cost of PoW mining) must be large relative to the one-off "stock" benefits of attacking it. In a previous work, we introduced zero-knowledge proof-of-identity (zk-PoI, [4]) for biometric passports [28] and electronic identity cards on permissionless blockchains in order to remove the inefficiencies of Sybil-resistance mechanisms such as proof-of-work [30] and proof-of-stake [37]: now, Sybil resistance only costs $O(1)$ instead of the PoW time complexity of $O(n \cdot b)$ for b blocks and n nodes as in Bitcoin. Additionally, attacks [38,39] on PoW sharded permissionless blockchains are prevented with zk-PoI: an identity is the same on all the shards, and the attacker cannot mine new identities for different shards as is possible on PoW blockchains.

Algorithm 1 Drand's rotating leader pseudo-code

```
NewViewChangeRequest(oldRoster) {
 nextRoster, time, round = rotateRosterRandomly(oldRoster)
 return NewViewReq(nextRoster, time, round)
}
rotateRosterRandomly(roster) {
 t = time(now)
 round = drand.RoundAt(t)
 rnd = drand.Get(round)
 mod = mod(rnd, length(roster))
 return newRoster(roster.List[mod:], roster.List[:mod]), t, round
}
ViewChangeVerifyRandom(block, request) {
 nextRoster = rotateRosterRandomVerification(block.Roster, request.time, request.round)
 if (nextRoster != request.Roster) error("invalid roster")
 if (!request.VerifyUniqueness()) error("proofs are not unique")
 threshold = defaultThreshold(length(block.Roster.List)
 numsigners = len(request.Proof)
 if (numsigners < threshold) error("not enough proofs")
 if (!request.VerifySignatures(block.Roster)) error("couldn't verify signatures")
}
rotateRosterRandomVerification(roster, time, roundBlock) {
 roundDrand = drand.RoundAt(time)
 if (roundDrand != roundBlock) error("different rounds")
 rnd = drand.Get(round)
 mod = mod(rnd, length(roster))
 return newRoster(roster.List[mod:], roster.List[:mod])
}
```

Although some could consider the latest zero-knowledge implementations to be fast enough, their implementations are still too experimental for production. For the first release, we will use the Software Guard Extensions (SGX) implementation based on mutual attestation, which works as described in the following Figure 1 (more details are in the original paper [4]):

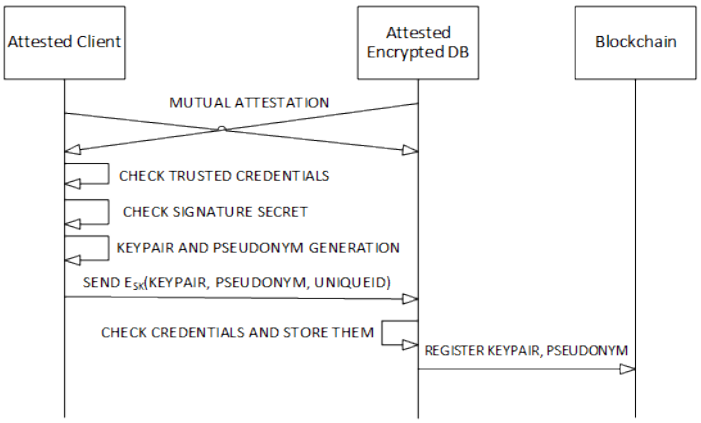

Figure 1. Simplified overview of mutual attestation protocol.

An approximate picture of the worldwide coverage follows in the next Figure 2:

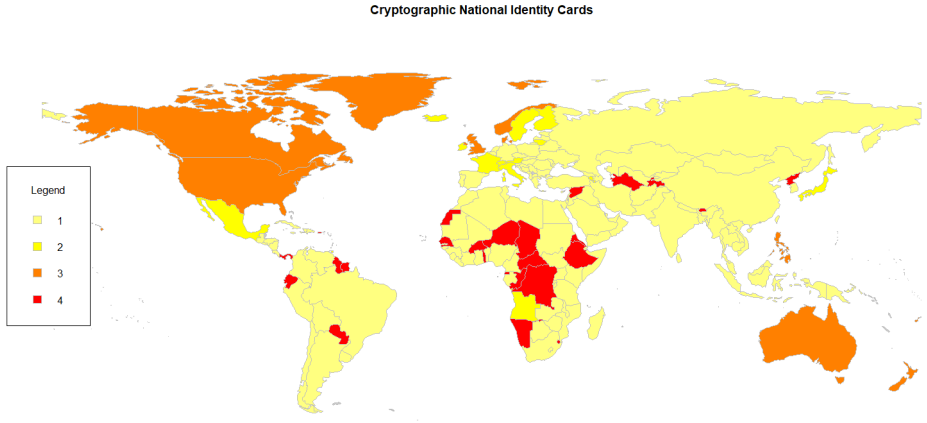

Figure 2. Legend: (1) national identity card is a mandatory smartcard; (2) national identity card is a voluntary smartcard; (3) no national identity card, but cryptographic identification is possible using an ePassport, drivers license, and/or health card; and (4) non-digital identity card.

For updates to worldwide coverage, check the URL of [40].

4.4. Implementation

Pravuil has a working implementation consisting of:

- A blockchain layer in Go and Java, invoking drand [31] as described in Section 4.2;
- Zero-knowledge proof-of-identity [4] in Python and C;
- Secure smart contracts in Obliv-Java [41].

5. Discussion

In this section, we discuss the economic rationale underpinning the unique features of this blockchain design that help it to overcome previous shortcomings and achieve an improved blockchain tailored to real-world settings according to the experiences from the last decade (e.g., Bitcoin [30]).

5.1. Overcoming Bitcoin's Limited Adoption Problem

In a recent paper [9], research showed that a PoW blockchain (e.g., Bitcoin) cannot simultaneously sustain a large volume of transactions and a non-negligible market share:

Proposition 1 (Adoption problem [9]). *Adoption decreases as demand rises (i.e., the adoption rate of a network, c^*, decreases in N). Moreover, the blockchain faces limited adoption,*

$$\lim_{N\to\infty} c^* = 0.$$

Even allowing dynamic PoW supply (i.e., by relaxing PoW artificial supply constraint) achieves widespread adoption only at the expense of decentralisation:

Proposition 2 ((Decentralisation implies limited adoption [9]). *PoW blockchains necessarily face either centralisation,*

$$\lim_{N\to\infty} \sup V \leq 1,$$

or limited adoption,

$$\lim_{N\to\infty} c^* = 0.$$

The previous propositions expose that the lack of widespread adoption constitutes an intrinsic property of PoW payment blockchains: as transaction demand grows, fees increase endogenously. Attracted by this growth, more nodes join the validation process, expanding the network size and thus protracting the consensus process and generating increased payment confirmation times: only users insensitive to wait times would transact in equilibrium, and limited adoption arises. Moreover, this limitation cannot be overcome, as it is rooted in physics (e.g., network delay).

As pointed out by the previous proposition, centralised blockchains overcome the limited adoption problem; for example, permissioned blockchains that remain secure on an open, adversarial network, such as the blockchain proposed in this paper, enable lower payment confirmation times when omitting PoW artificial supply constraint.

Proposition 3 (Lower payment confirmation times [9]). *For any PoW protocol, there exists a permissioned blockchain that remains secure on an open, adversarial network (e.g., Pravuil), which induces (weakly) lower payment confirmation times.*

Additionally, omitting PoW artificial supply constraint facilitates timely service even for high transaction volumes:

Proposition 4 (No limited adoption problem [9]). *In any permissioned equilibrium, widespread adoption can be obtained,*

$$\lim_{N\to\infty} c_P^* = \min\left\{\frac{e_P}{\Delta(V_P)}, 1\right\}.$$

Proofs for all the previous propositions can be found in the corresponding paper.

5.2. Obtaining Higher Transaction Security at a Lower Cost

In another recent paper [8], the authors show that permissioned blockchains have a higher level of transaction safety than permissionless blockchains, independent of the block reward and the current exchange rate of the cryptocurrency.

For a PoW permissionless blockchain, let R be the block reward in the corresponding cryptocurrency, x the associated exchange rate to fiat currency, w the block maturation rate (e.g., for Bitcoin, $R = 6.25$; x = \$60,000; w = 100), f the probability of detecting that blocks have been replaced, and β_{pl} the value above which transactions are not safe,

$$\beta_{pl} = fwRx.$$

Note that 51% attacks are becoming more common, specially for purely financial reasons [42]. For a permissioned blockchain, let p_i be the punishment applied to each node i if it participates in an attack, $\tau \in [0, 1]$ be the probability that nodes that participated in an attack are punished, and β_P be the value above which transactions are not safe,

$$\beta_P = f\tau \sum_{i\in B} p_i,$$

with B being the set of n nodes with the lowest p_i. Typical punishments include confiscating all the funds deposited on the blockchain and banning the user(s) from the blockchain, among others

Proposition 5 ([8]). *A permissioned blockchain that is safe in an open, adversarial environment (e.g., Pravuil) has a higher level of maximum value for transaction safety than a PoW permissionless blockchain if*

$$\tau \sum_{i\in B} p_i > wRx.$$

Even small values of τ result in higher safety for larger transactions than PoW permissionless blockchains:

Proposition 6 ([8]). *For $\tau > 0$ and high enough p_is, a permissioned blockchain that is safe in an open, adversarial environment (e.g., Pravuil) is more resilient than PoW permissionless blockchains whenever*

$$\sum_{i \in B} p_i > \frac{wRx}{\tau}.$$

Ultimately, the cost of providing incentives to the validators so that they do not participate in potential attacks (e.g., validator incentives such as block rewards) will be lower for permissioned blockchains.

Proofs for all the previous propositions can be found in the corresponding paper.

Proposition 7 ([8]). *Suppose that $\beta_{pl} > 0$ and $\beta_P > 0$; then, at equilibrium, the validator incentives in the permissioned blockchain that is safe in an open, adversarial environment (e.g., Pravuil) are lower than those for PoW permissionless blockchains.*

According to the model of this paper, in order to increase transaction safety, we only need to increase:

- τ, a probability that reflects user's trust in the system;
- p_i, a penalty that could also include legal action.

In general, the mere existence of credible penalties p_i with positive probability τ is enough for the system to remain secure without needing to exert punishments in the case of rational attackers. Additionally, note that these parameters are not economic parameters of the system, unlike the parameters for PoW permissionless blockchains.

5.3. An Empirical Approach to Blockchain Design

Motivated by the abstract analysis from the previous Section 5.2, we use the numerical comparisons between cryptocurrencies from [6] to compare permissionless and permissioned blockchains in practice.

The data in the previous Tables 2 and 3 were obtained via a focus group through semi-structured interviews with blockchain practitioners and researchers in order to come up with numerical assessments (i.e., scores in the range [1, 5]) of the top cryptocurrencies at the moment it was conducted; thus, the popularity scores differ from the current market capitalisation. For the sake of completeness, all the different features are hereby described: regarding cost, the lowest scores correspond to higher transaction fees (i.e., excluding Sybil-resistance mechanism); for consistency, lower scores indicate higher confirmation times while higher scores indicate fast consistency, usually obtained with newer consensus protocols; for functionality, higher scores denote the support of smart contracts; for performance, higher scores correspond to the fastest block production rates and transactions per second; for security, higher scores indicate that the adversary needs to control a higher proportion of malicious mining nodes; for decentralisation, lower scores denote more centralised blockchains while higher scores denote decentralised blockchains without trusted authorities nor hierarchies among the participants; for privacy, lower scores indicate blockchains without privacy solutions to obfuscate transactions and/or identities.

Table 2. Permissionless blockchains. (*): $p < 0.05$.

Features	BTC	ETH	BCH	LTC	ADA	USDT	Mean
Popularity	1	2	3	7	8	9	
Cost	1.33	2	1.66	2.66	4.33	5	2.83 (*)
Consistency	1.33	2.33	1.33	2	3.66	1	
Functionality	2	5	2	2	4.33	2	
Performance	1.33	1.66	2	2.33	3	1	1.88 (*)
Security	4	4	4	4	4	3.33	3.88
Decentralisation	5	3.33	4.33	3.66	3.33	1.33	
Total	14.99	18.32	15.32	16.7	22.65	13.66	
Performance/Cost	0.28	0.41	0.46	0.7	1.79	1	0.77 (*)
(Perf*Sec)/Cost	1.13	1.66	1.84	2.79	7.18	3.33	2.99 (*)
Security/Cost	0.85	1	0.92	1.2	2.39	3.33	1.61 (*)

Table 3. Permissioned blockchains. (*): $p < 0.05$.

Features	XRP	EOS	XLM	TRX	MIOTA	Mean
Popularity	3	5	6	11	10	
Cost	4.66	5	4.66	5	5	4.84 (*)
Consistency	4.33	5	4	4	4.66	
Functionality	1.33	5	1.33	5	3.66	4.53 (*)
Performance	4.33	4.66	4	4.66	5	3.33
Security	2.33	3.33	4	3.33	3.66	
Decentralisation	1	2.66	2.33	3.33	2.33	
Total	17.98	25.65	20.32	25.32	24.31	
Performance/Cost	3.23	4.66	2.98	4.66	5	4.10 (*)
(Perf*Sec)/Cost	7.52	15.51	11.94	15.51	18.3	13.76 (*)
Security/Cost	1.73	3.33	2.98	3.33	3.66	3 (*)

Using two-sample *t*-tests assuming unequal variances, we compare the following means between permissionless and permissioned blockchains, remarking that they are statistically significant:

- Cost: permissionless blockchains are costlier (2.83) than permissioned blockchains (4.84). Please note that a higher cost score means that the blockchain is considered to have better costs (i.e., lower costs), and the ranking obtained from this cost score must be reversed to be useful in the next rankings.
- Performance: permissionless blockchains are less performant (1.88) than permissioned blockchains (4.53).
- Performance/Cost: permissionless blockchains show worse performance regarding cost (0.77) than permissioned blockchains (4.10).
- (Performance*Security)/Cost: permissionless blockchains show worse performance and security regarding cost (2.99) than permissioned blockchains (13.76).
- Security/Cost: permissionless blockchains show worse security regarding cost (1.61) than permissioned blockchains (3).

It is clear from the empirical data that permissionless blockchains are considered worse than permissioned blockchains when considering cost, performance, and security.

5.4. Achieving More Decentralisation than Other Permissionless Blockchains

In yet another recent publication [7], the writers noticed that permissioned blockchains could achieve more decentralisation than permissionless blockchains: real-world permissionless blockchains are quite centralised [43] as there are not formal checks for the underlying centralisation.

In order to obtain a more decentralised permissioned blockchain that admission/gatekeeping function must be decentralised and opened; precisely, is safe in an open, adversarial network (e.g., Pravuil in next Figure 3), the node ideal state is achieved with zero-knowledge proof-of-identity [4], as previously explained in Section 4.3.

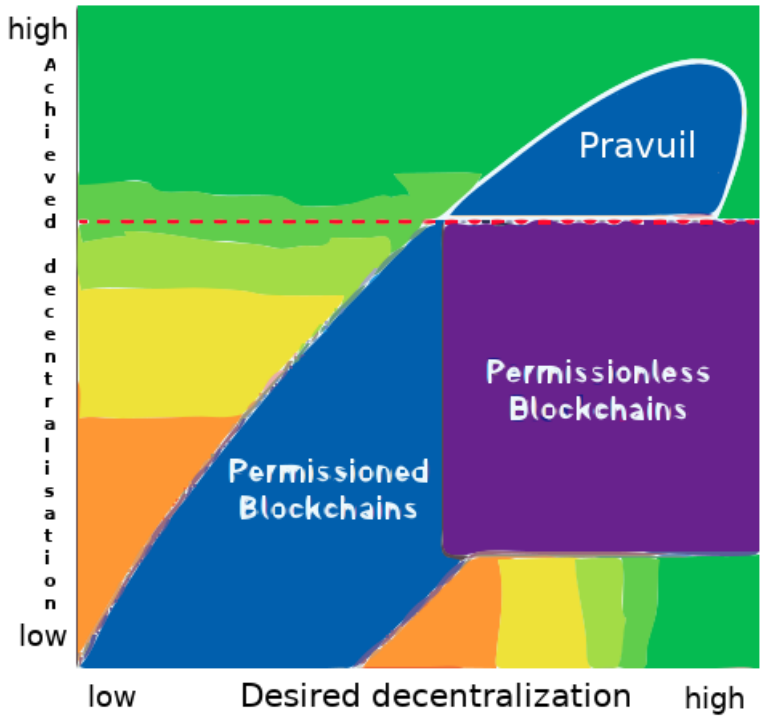

Figure 3. Comparing decentralisation (from [7]).

5.5. Overcoming the Scalability Trilemma

The scalability trilemma postulates that a blockchain system can at most have only two of the following three properties: decentralisation, scalability, and security.

In Pravuil, decentralisation, scalability, and security can be achieved simultaneously as pictured in Figure 4:

- Decentralisation: as discussed previously in Section 5.4, Pravuil can be more decentralised than other permissionless blockchains by using zero-knowledge proof-of-identity, as previously explained in Section 4.3. It also circumvents the impossibility of full decentralisation [4].
- Scalability: Pravuil inherits the scalable *rotating-subleader* (RS) communication pattern from MOTOR [19], specifically created to avoid the communication bottleneck experienced by classic BFT protocols when run over limited-bandwidth networks.
- Security: Pravuil is secure, as previously proven in Theorems 1 and 2.

Note that overcoming the scalability trilemma does not imply that the consistency–availability–partition theorem is also overcome: in case of network partition, consistency is ensured over availability (i.e., because the consensus is strongly consistent, as previously defined).

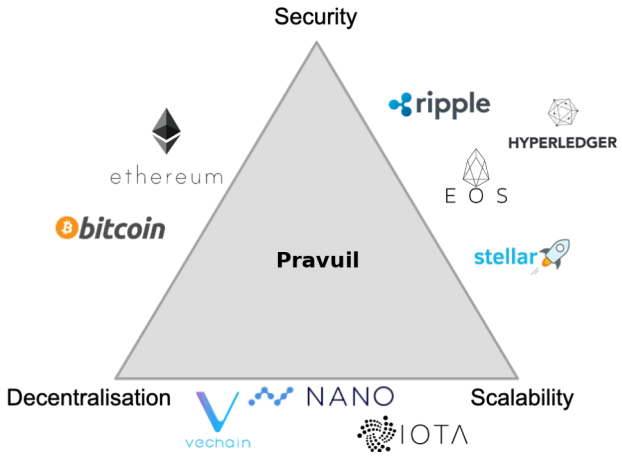

Figure 4. Pravuil overcomes the scalability trilemma.

5.6. Obviating the Price of Crypto-Anarchy of PoW/PoS Cryptocurrencies

In a previous paper [4], we pointed out that the most cost-efficient Sybil-resistance mechanism is the one provided by a trusted national PKI infrastructure [44], and a centralised social planner would prefer the use of national identity cards and/or ePassports in order to minimise costs: instead, permissionless blockchains are paying very high costs by using PoW/PoS as Sybil-resistance mechanisms. The Price of Crypto-Anarchy compares the ratio between the worst Nash equilibrium of the congestion game defined by PoW blockchains and the optimal centralised solution, quantifying the costs of the selfish behaviour of miners.

Definition 9 (#26 from [4]). *Let $NashCongestedEquil \subseteq S$ be the set of strategies given as the solution to the optimisation problem of Theorem 25 from [4]; then, the Price of Crypto-Anarchy is given by the following ratio:*

$$Price\ of\ Crypto\text{-}Anarchy = \frac{\max_{s \in NashCongestedEquil} Cost(s)}{Cost(zk\text{-}PoI)}$$

In practice, the real-world costs of zero-knowledge proof-of-identity are almost zero, as the identity infrastructure is subsidised by governments. However, the situation for PoW/PoS blockchains is quite the opposite:

- PoW blockchains: in 2018, Bitcoin, Ethereum, Litecoin, and Monero consumed an average of 17, 7, 7, and 14 MJ, respectively, to generate one USD [45], and in 2021, Bitcoin may have consumed as much energy as all data centres globally [46,47] at 100–130 TWh per year. Holders of cryptocurrency ultimately experience the Price of Crypto-Anarchy as inflation from mining rewards; the next Table 4 compares the yearly mining rewards of some of the top PoW cryptocurrencies and provides estimations of their yearly inflation rates; note that for Ether, the table is only for PoW and before the merge transitioned Ethereum to proof-of-stake (PoS), because although the inflation rate is much lower with PoS (even deflationary), there was not a long enough time series at the time of publication to provide a reliable estimate of the inflation rate.

Table 4. Empirical Price of Crypto-Anarchy.

Name	Reward per Block	Block Time	Blocks per Day	Price	Yearly Mining Reward	Yearly Inflation
BTC	6.25	10 m	144	$50,000	$18,061 B	4.12%
ETH	2	13.2 s	6545	$3780	$16,425 B	1.76%
DOGE	10,000	1 m	1440	$0.49	$2575 B	4.06%
LTC	12.5	2.5 m	576	$320	$840 MM	3.94%
BCH	6.25	10 m	144	$1275	$418 MM	1.75%
ZEC	3.125	75 s	1152	$301	$395 MM	11.84%
XMR	1.02	2 m	720	$407	$109 MM	1.5%

- PoS blockchains: in theory, the costs are identical to the costs of PoW schemes, except that instead of electrical resources and mining chips, they take the form of illiquid financial resources [11], and in practice, proof-of-stake is not strictly better than proof-of-work, as the distribution of the market shares between both technologies has been shown to be indistinguishable (Appendix 3, [48]).

Bitcoin miners have earned a total of $26.75B as of April 2021. It is not necessary to pay so much for Sybil resistance; instead, miners could be paid for other tasks (e.g., transaction fees). As previously discussed, obtaining Sybil resistance for free is not only the key to overcoming Bitcoin's limited adoption problem (Section 5.1) and achieving more decentralisation than other permissionless blockchains (Section 5.4), but also to going beyond the economic limits of Bitcoin, as discussed in the next Section 5.7.

5.7. Beyond the Economic Limits of Bitcoin

In a paper about the economic limits of Bitcoin [10], the author pointed out that Bitcoin is prohibitively expensive to run because the recurring "flow" payments to miners for running the blockchain (particularly, the cost of PoW mining) must be large relative to the one-off "stock" benefits of attacking it. Let V_{attack} be the expected payoff to the attacker, R_{block} be the block reward to the miner, and α represent the duration of the attack net of block rewards; then,

$$R_{block} > \frac{v_{attack}}{\alpha},$$

placing serious economic constraints on the practicality and scalability of the Bitcoin blockchain, a problem that seems intrinsic to any anonymous, decentralised blockchain protocol. Consequently, the author poses the open question of finding another approach to generating anonymous, decentralised trust in a public ledger that is less economically expensive; indeed, the technical solution presented in Section 4.3 that incorporates zero-knowledge proof-of-identity [4] is the technology that is both "scarce and non-repurposable", affordable and not susceptible to sabotage attacks that could cause a collapse in the economic value of the blockchain that the author of [10] would seem meritorious to close said open question.

A more recent paper [11] continues the previous economic analysis [11], extending it to PoS and permissioned settings. For the permissionless PoS setting, the authors find that the costs are identical to the cost of PoW schemes, except that instead of electrical resources and mining chips, they takes the form of illiquid financial resources; however, zk-PoI [4] is free. For the permissioned case concerning this paper, if the block reward is set exogenously, a permissioned blockchain has lower costs than permissionless PoW or PoS blockchains in the economic model of [10].

5.8. More Valuable and Stable Cryptocurrencies

A review of previous literature in economic research reveals the following interesting facts regarding the intricate relationship between PoW mining (e.g., hashrate, electricity, and/or equipment costs) and cryptocurrency prices:

- There is a positive relationship between mining hashrate and price [49,50]; the causality is primarily unidirectional, going from the price to the hashrate [12], although mining incidents and political shocks that affect mining also negatively impact prices.
- Bitcoin's security is sensitive (elastic) to mining rewards and costs, although temporary mining costs and price shocks do not affect the long-run blockchain security [13]; a 1% permanent increase in the mining reward increases the underlying blockchain security by 1.38% to 1.85% in the long-run; positive shocks to electricity prices in China have a negative impact on the hashrate in the short-run; a 1% increase in the efficiency of mining equipment increases the computing capacity between 0.23% and 0.83% in the long-run; in the short-run, mining competition intensity has a statistically positive impact, leading to expansion of mining capacity, but in the long-run, the relationship is reversed.
- High fixed mining rewards are the source of the instability in reaching an equilibrium between miners and users [16]; instead, mining rewards should be adjusted dynamically.
- The production of cryptocurrency by miners is jointly determined by the price used by consumers [15]; the equilibrium price depends on both consumer preferences (i.e., price increases with the average value of censorship aversion and the current and future size of the network) and industrial organisation of the mining market (i.e., price increases with the number of miners and decreases with the marginal cost of mining). Price-security spirals amplify demand and supply shocks: for example, a sudden demand shock provoked by a government banning the cryptocurrency in a country would lead to price drops, leading to miners decreasing hashrate, further decreasing prices, and the feedback loop continues until a new equilibrium is reached in multiple rounds. In other words, Bitcoin's security model embeds price volatility amplification: indeed, empirical studies [51] find evidence of persistence and long memory volatility in Bitcoin's market.
- In a PoW blockchain, it is impossible to simultaneously achieve all three following goals [14]: maximise cryptocurrency price, blockchain security, and social welfare.

Similar results can be found for PoS blockchains because they substitute electricity and mining costs for illiquid and volatile financial resources [11]. In general, the interdependencies can be described graphically as the following cycles in Figure 5 and spirals in Figure 6:

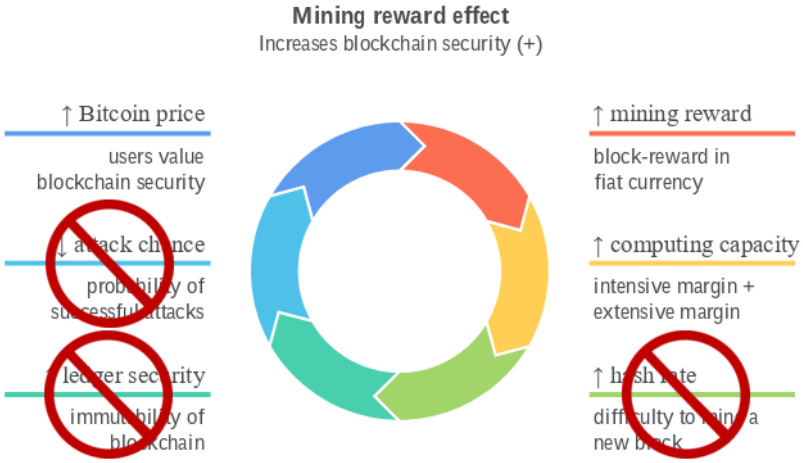

Figure 5. Interdependencies [13] with broken negative feedback loops.

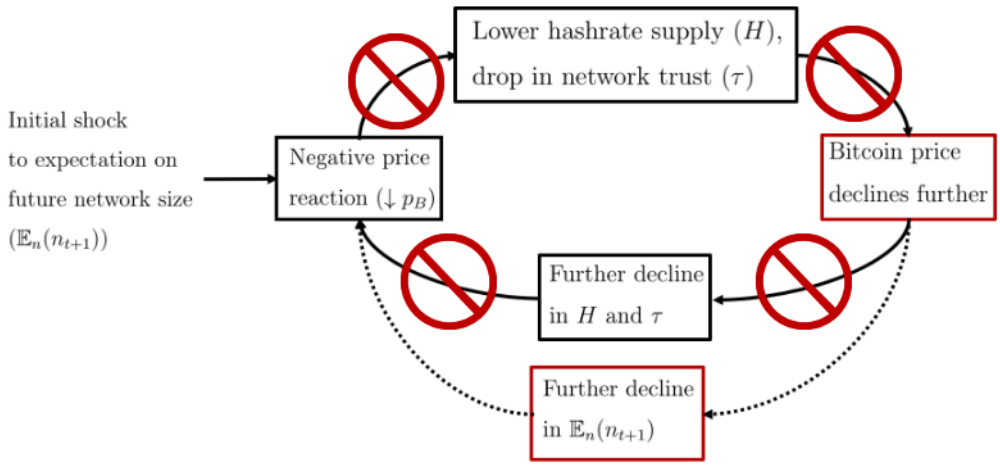

Figure 6. Spirals [15] with broken negative feedback loops.

However, we break most of the previous interdependencies and spirals with our strongly consistent blockchain with free Sybil resistance:

- Blockchain and transaction security are independent of blockchain mining capacity, mining costs and rewards, and price: once a transaction is instantly committed, it is committed forever;
- There are not price–security spirals for demand and supply shocks: changes in prices do not lead to changes in security;
- As blockchain security is independent of price, it is possible to maximise cryptocurrency price and social welfare.

Ultimately, our blockchain design leads to more valuable and stable cryptocurrencies.

6. Evaluation

An open-source implementation of Pravuil is available at (Supplementary materials, 2 September 2022) in multiple repositories. Pravuil's random leader rotation is implemented by interrogating the drand service on each round in order to obtain a public, unpredictable, unbiasable source of randomness that safely elects the next leader: its implementation is quite straightforward, as the most complex part, the randomness source, is outsourced to drand. Pravuil also adds zero-knowledge proof-of-identity (zk-PoI) by creating a new service on the blockchain that listens to new registration requests coming from only zk-PoI's SGX implementation, thus making its architecture more modular and improving its high availability by enabling the use of another dedicated set of servers as registration nodes.

To better understand the scalability of Pravuil under failures, we simulate the consensus protocol between mining nodes on a cloud deployment with 20 hosts and 50 Pravuil nodes per host; additionally, the network environment is configured with an RTT of 180 ms and a maximum bandwidth of 25 Mbps. Note that the following results may vary under different network configurations and different transaction loads (e.g., simple transactions vs. complex smart contracts).

In all the following figures, consensus group size denotes the number of nodes in only one shard (as shown in the graphs below, the maximum recommended size is 600 nodes), and latency denotes the time to commit a transaction on the blockchain, calculated as an average of 200 transactions: lower latency is better for all figures.

In Figure 7, Pravuil always shows much better latency than its predecessor (e.g., ByzCoin), irrespective of the number of nodes and block size (10 bytes vs. 1 MB, Bitcoin's block size).

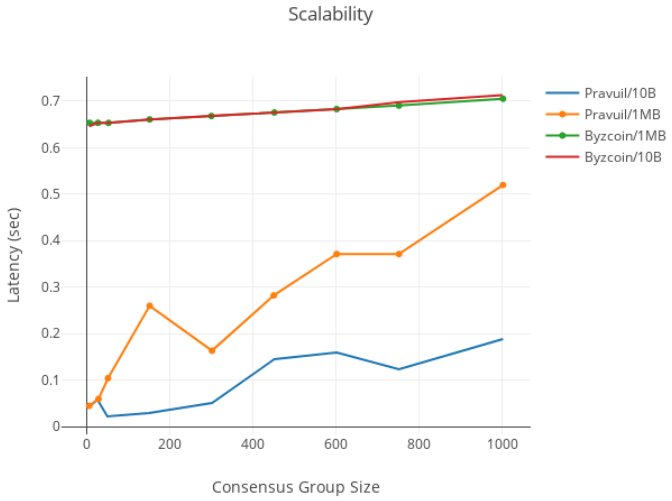

Figure 7. Scalability.

In Figure 8, Pravuil shows its fault tolerance under different configurations of number of nodes: large scale deployments of hundreds of nodes still reach consensus in a few seconds even with a high number of failures.

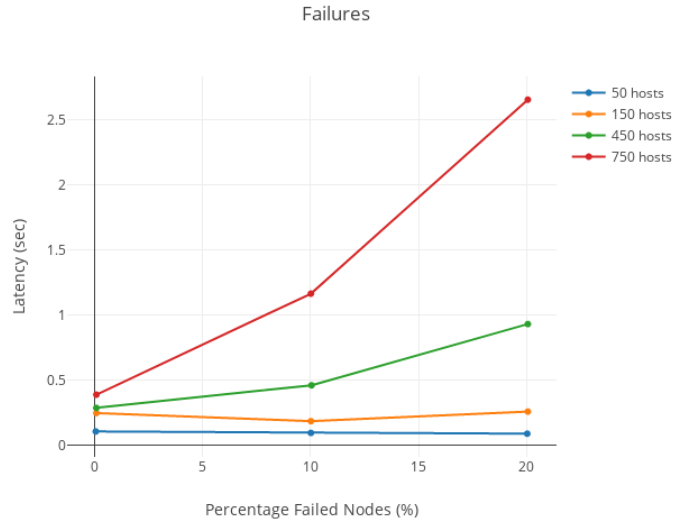

Figure 8. Failures.

In Figure 9, the scalability of Pravuil without faults (0 percent) is shown to be quite consistent, but with faults (20 percent of nodes failing), it gets only slightly worse (always under 1 s) as long as the number of nodes is less than 600 (i.e., the maximum recommended number of nodes for a consensus group).

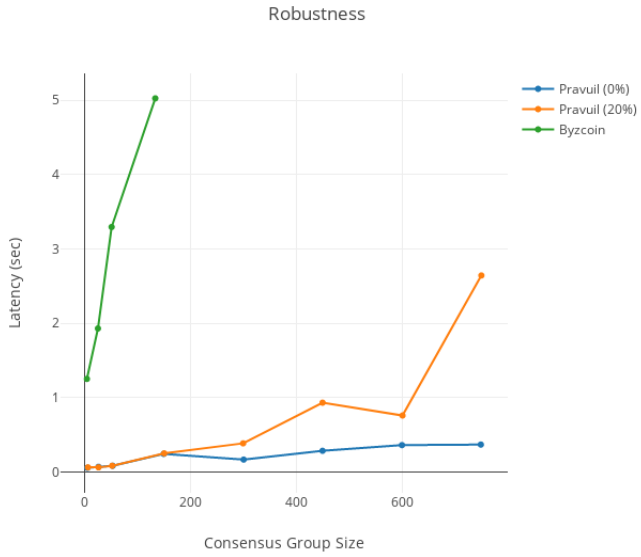

Figure 9. Robustness.

7. Conclusions

In this paper, following the research design set forth in Section 3.2, a broad literature review was conducted about the economic consequences of the technical features used by blockchains and cryptocurrencies (Section 5), leading to the optimal architecture described in Section 4, featuring:

- Unpredictably rotating leaders using drand to defend against adversaries and censorship attacks;
- Zero-knowledge proof-of-identity [4] as a Sybil-resistance mechanism to overcome Bitcoin's limited adoption problem [9] and to go beyond the economic limits of Bitcoin [10], delivering more decentralisation than other permissionless blockchains [7].

An implementation (Section 4.4) and simulations under different conditions (Section 6 concluded that Pravuil is certainly an improvement with better performance over previous blockchains, is suitable for real-world deployment in adversarial networks such as the Internet, and features much better economic properties as was sought in Section 5.

Further research could consider the interrelationship between the reliability of the blockchain, its functional safety, and economic indicators: for example, within the recently introduced framework of [52]. Additionally, further research should address the obvious limitation of the best way to handle rogue states when, in the future, Pravuil will be running in production.

Supplementary Materials: Code being open-sourced and ongoing development is available at (2 September 2022): https://github.com/Calctopia-OpenSource.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BFT	Byzantine fault tolerance
BLS	Boneh–Lynn–Sacham
DKG	distributed key generation
FATF	Financial Action Task Force
GST	global stabilisation time
PKI	public key infrastructure
PoI	proof-of-identity
PoS	proof-of-stake
PoW	proof-of-work
SGX	Software Guard eXtensions
zk-PoI	zero-knowledge proof-of-identity

Notations

The following notations are used in this manuscript:

Variable	Definition
m	message
P_i	party i
Δ	network time bound
n	number of nodes
f	faulty nodes
d	attacker-controlled nodes
c^*	network adoption rate
N	number of users
V	transaction rate
e	exogenous parameter
R	block reward
x	crypto-to-fiat exchange rate
w	block maturation rate
f	probability of detecting block replacement
β_{pl}	permissionless transaction safety upper bound
β_P	permissioned transaction safety upper bound
p_i	punishment for node i
τ	probability of punishing attacker's node
V_{attack}	expected attacker's payoff
α	duration of the attack
$=$	assignment
$\in$	element of
$\subseteq$	subset of

References

1. Garay, J.; Kiayias, A. SoK: A Consensus Taxonomy in the Blockchain Era. In *Topics in Cryptology–CT-RSA 2020*; Jarecki, S., Ed.; Springer International Publishing: Cham, Switzerland, 2020; pp. 284–318. [CrossRef]
2. Liu, Y.; Liu, J.; Salles, M.A.V.; Zhang, Z.; Li, T.; Hu, B.; Henglein, F.; Lu, R. Building Blocks of Sharding Blockchain Systems: Concepts, Approaches, and Open Problems. *arXiv* **2021**, arXiv:2102.13364.
3. Raikwar, M.; Gligoroski, D.; Kralevska, K. SoK of Used Cryptography in Blockchain. *IEEE Access* **2019**, *7*, 148550–148575. [CrossRef]
4. Cerezo Sánchez, D. Zero-Knowledge Proof-of-Identity: Sybil-Resistant, Anonymous Authentication on Permissionless Blockchains and Incentive Compatible, Strictly Dominant Cryptocurrencies. Cryptology ePrint Archive, Report 2019/546. 2019. Available online: https://eprint.iacr.org/2019/546 (accessed on 2 September 2022).
5. Bonneau, J.; Miller, A.; Clark, J.; Narayanan, A.; Kroll, J.A.; Felten, E.W. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In Proceedings of the 2015 IEEE Symposium on Security and Privacy, San Jose, CA, USA, 17–21 May 2015; pp. 104–121. [CrossRef]
6. Garriga, M.; Dalla Palma, S.; Arias, M.; De Renzis, A.; Pareschi, R.; Andrew Tamburri, D. Blockchain and cryptocurrencies: A classification and comparison of architecture drivers. *Concurr. Comput. Pract. Exp.* **2021**, *33*, e5992. [CrossRef]

7. Bakos, Y.; Halaburda, H.; Mueller-Bloch, C. When Permissioned Blockchains Deliver More Decentralization Than Permissionless. *Commun. ACM* **2021**, *64*, 20–22. [CrossRef]
8. Bakos, Y.; Halaburda, H. Tradeoffs in Permissioned vs Permissionless Blockchains: Trust and Performance. 2021. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3789425 (accessed on 2 September 2022).
9. Hinzen, F.J.; John, K.; Saleh, F. Bitcoin's limited adoption problem. *J. Financ. Econ.* **2022**, *144*, 347–369. [CrossRef]
10. Budish, E. The Economic Limits of Bitcoin and the Blockchain. 2018. Available online: https://www.nber.org/system/files/working_papers/w24717/w24717.pdf (accessed on 2 September 2022).
11. Gans, J.S.; Gandal, N. More (or Less) Economic Limits of the Blockchain. 2019. Available online: https://www.nber.org/system/files/working_papers/w26534/w26534.pdf (accessed on 2 September 2022).
12. Fantazzini, D.; Kolodin, N. Does the Hashrate Affect the Bitcoin Price? *J. Risk Financ. Manag.* **2020**, *13*, 263. [CrossRef]
13. Ciaian, P.; d'Artis Kancs.; Rajcaniova, M. Interdependencies between Mining Costs, Mining Rewards and Blockchain Security. *Ann. Econ. Financ.* **2021**, *22*, 1–36.
14. Pagnotta, E.S. Decentralizing Money: Bitcoin Prices and Blockchain Security. *Rev. Financ. Stud.* **2021**, *35*, 866–907. [CrossRef]
15. Pagnotta, E.; Buraschi, A. An Equilibrium Valuation of Bitcoin and Decentralized Network Assets. 2018. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3142022 (accessed on 2 September 2022).
16. Iyidogan, E. An Equilibrium Model of Blockchain-Based Cryptocurrencies. 2018. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3152803 (accessed on 2 September 2022).
17. 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 the 25th USENIX Conference on Security Symposium, Anaheim, CA, USA, 10–12 August 2016; USENIX Association: Austin, TX, USA, 2016; pp. 279–296. [CrossRef]
18. 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 the 2018 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 20–24 May 2018; pp. 583–598. [CrossRef]
19. Kokoris-Kogias, E. Robust and Scalable Consensus for Sharded Distributed Ledgers. Cryptology ePrint Archive, Report 2019/676. 2019. Available online: https://eprint.iacr.org/2019/676 (accessed on 2 September 2022).
20. Damgård, I.; Ganesh, C.; Khoshakhlagh, H.; Orlandi, C.; Siniscalchi, L. Balancing Privacy and Accountability in Blockchain Identity Management. In *Topics in Cryptology—CT-RSA 2021*; Paterson, K.G., Ed.; Springer International Publishing: Cham, Switzerlan, 2021; pp. 552–576. [CrossRef]
21. Ramasamy, H.V.; Cachin, C. Parsimonious Asynchronous Byzantine-Fault-Tolerant Atomic Broadcast. In *Principles of Distributed Systems*; Anderson, J.H., Prencipe, G., Wattenhofer, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2006; pp. 88–102. [CrossRef]
22. Dolev, D.; Dwork, C.; Stockmeyer, L. On the Minimal Synchronism Needed for Distributed Consensus. *J. ACM* **1987**, *34*, 77–97. [CrossRef]
23. Dwork, C.; Lynch, N.; Stockmeyer, L. Consensus in the Presence of Partial Synchrony. *J. ACM* **1988**, *35*, 288–323. [CrossRef]
24. Fischer, M.J.; Lynch, N.A.; Merritt, M. Easy Impossibility Proofs for Distributed Consensus Problems. In Proceedings of the Fourth Annual ACM Symposium on Principles of Distributed Computing, Minaki, ON, Canada, 5–7 August 1985; Association for Computing Machinery: New York, NY, USA, 1985; pp. 59–70. [CrossRef]
25. Pass, R.; Shi, E. Hybrid Consensus: Efficient Consensus in the Permissionless Model. In Proceedings of the 31st International Symposium on Distributed Computing (DISC 2017), Vienna, Austria, 16–20 October 2017; Richa, A.W., Ed.; Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik: Dagstuhl, Germany, 2017; Volume 91, Leibniz International Proceedings in Informatics (LIPIcs), pp. 39:1–39:16. [CrossRef]
26. Boneh, D.; Lynn, B.; Shacham, H. Short Signatures from the Weil Pairing. In *Advances in Cryptology—ASIACRYPT 2001*; Boyd, C., Ed.; Springer: Berlin/Heidelberg, Germany 2001; pp. 514–532. [CrossRef]
27. Boneh, D.; Drijvers, M.; Neven, G. Compact Multi-signatures for Smaller Blockchains. In *Advances in Cryptology—ASIACRYPT 2018*; Peyrin, T., Galbraith, S., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 435–464. [CrossRef]
28. ICAO. ICAO Public Key Directory (PKD). 2022. Available online: https://www.icao.int/Security/FAL/PKD/Pages/default.aspx (accessed on 2 September 2022).
29. Stifter, N.; Judmayer, A.; Schindler, P.; Kern, A.; Fdhila, W. What Is Meant by Permissionless Blockchains? Cryptology ePrint Archive, Report 2021/023. 2021. Available online: https://eprint.iacr.org/2021/023 (accessed on 2 September 2022).
30. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. 2009. Available online: https://bitcoin.org/bitcoin.pdf (accessed on 2 September 2022).
31. DRAND. DRAND. 2021. Available online: https://drand.love (accessed on 2 September 2022).
32. DRAND. DRAND Specification. 2021. Available online: https://drand.love/docs/specification (accessed on 2 September 2022).
33. Feldman, P. A Practical Scheme for Non-Interactive Verifiable Secret Sharing. In Proceedings of the 28th Annual Symposium on Foundations of Computer Science, Los Angeles, CA, USA, 12–14 October 1987; IEEE Computer Society: New York, NY, USA, 1987; pp. 427–438. [CrossRef]
34. Gennaro, R.; Jarecki, S.; Krawczyk, H.; Rabin, T. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. In *Advances in Cryptology—EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques*; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 1999; Volume 1592, pp. 295–310. [CrossRef]

35. Galindo, D.; Liu, J.; Ordean, M.; Wong, J.M. Fully Distributed Verifiable Random Functions and their Application to Decentralised Random Beacons. In Proceedings of the 2021 IEEE European Symposium on Security and Privacy (EuroSP), Vienna, Austria, 6–10 September 2021; pp. 88–102. [CrossRef]
36. DRAND. DRAND Security Model. 2021. Available online: https://drand.love/docs/security-model (accessed on 2 September 2022).
37. King, S.; Nadal, S. PPCoin: Peer-to-Peer Crypto-Currency with Proof-of-Stake. 2012. Available online: https://www.peercoin.net/whitepapers/peercoin-paper.pdf (accessed on 2 September 2022).
38. Rajab, T.; Manshaei, M.H.; Dakhilalian, M.; Jadliwala, M.; Rahman, M.A. On the Feasibility of Sybil Attacks in Shard-Based Permissionless Blockchains. *arXiv* **2020**, arXiv:2002.06531.
39. Hafid, A.; Hafid, A.; Samih, M. A Tractable Probabilistic Approach to Analyze Sybil Attacks in Sharding-Based Blockchain Protocols. *IEEE Trans. Emerg. Top. Comput.* **2022**. [CrossRef]
40. Cerezo Sánchez, D. Updates to Worldwide Coverage. 2022. Available online: https://www.calctopia.com/category/coverage/ (accessed on 2 September 2022).
41. Cerezo Sánchez, D. Raziel: Private and Verifiable Smart Contracts on Blockchains. Cryptology ePrint Archive, Report 2017/878. 2017. Available online: https://eprint.iacr.org/2017/878 (accessed on 2 September 2022).
42. Shanaev, S.; Shuraeva, A.; Vasenin, M.; Kuznetsov, M. Cryptocurrency Value and 51% Attacks: Evidence from Event Studies. *J. Altern. Investments Winter* **2020**. [CrossRef]
43. Gencer, A.E.; Basu, S.; Eyal, I.; van Renesse, R.; Sirer, E.G. Decentralization in Bitcoin and Ethereum Networks. In *Financial Cryptography and Data Security*; Meiklejohn, S., Sako, K., Eds.; Springer: Berlin/Heidelberg, Germany, 2018; pp. 439–457. [CrossRef]
44. Douceur, J. The Sybil Attack. In Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS), Cambridge, MA, USA, 7–8 March 2002. [CrossRef]
45. Krause, M.J.; Tolaymat, T. Quantification of energy and carbon costs for mining cryptocurrencies. *Nat. Sustain.* **2018**, *1*, 711–718. [CrossRef]
46. Digiconomist. Bitcoin May Consume as Much Energy as All Data Centers Globally. 2021. Available online: https://digiconomist.net/bitcoin-may-consume-as-much-energy-as-all-data-centers-globally (accessed on 2 September 2022).
47. de Vries, A. Bitcoin Boom: What Rising Prices Mean for the Network's Energy Consumption. *Joule* **2021**, *5*, 509–513. [CrossRef]
48. ElBahrawy, A.; Alessandretti, L.; Kandler, A.; Pastor-Satorras, R.; Baronchelli, A. Evolutionary dynamics of the cryptocurrency market. *R. Soc. Open Sci.* **2017**, *4*, 170623. [CrossRef] [PubMed]
49. Georgoula, I.; Pournarakis, D.; Bilanakos, C.; Sotiropoulos, D.; Giaglis, G.M. Using Time-Series and Sentiment Analysis to Detect the Determinants of Bitcoin Prices. MCIS 2015 PROCEEDINGS. 2015. Available online: https://aisel.aisnet.org/mcis2015/20 (accessed on 2 September 2022).
50. Hayes, A.S. Cryptocurrency value formation: An empirical study leading to a cost of production model for valuing bitcoin. *Telemat. Inf.* **2016**, *34*, 1308–1321. [CrossRef]
51. Bouri, E.; Gil-Alana, L.A.; Gupta, R.; Roubaud, D. Modelling long memory volatility in the Bitcoin market: Evidence of persistence and structural breaks. *Int. J. Financ. Econ.* **2019**, *24*, 412–426. [CrossRef]
52. Kovtun, V.; Izonin, I.; Greguš, M. Model of Information System Communication in Aggressive Cyberspace: Reliability, Functional Safety, Economics. *IEEE Access* **2022**, *10*, 31494–31502. [CrossRef]