

# On the Security of Permissioned Blockchain Solutions for IoT Applications

Sotirios Brotsis\*, Nicholas Kolokotronis\*, Konstantinos Limniotis†, Stavros Shiaeles‡

\*University of Peloponnese, Greece. Email: {brotsis,nkolok}@uop.gr

†Hellenic Data Protection Authority, Greece. Email: klimniotis@dpa.gr

‡University of Portsmouth, UK. Email: stavros.shiaeles@port.ac.uk

**Abstract**—The blockchain has found numerous applications in many areas with the expectation to significantly enhance their security. The *Internet of things* (IoT) constitutes a prominent application domain of blockchain, with a number of architectures having been proposed for improving not only security but also properties like transparency and auditability. However, many blockchain solutions suffer from inherent constraints associated with the consensus protocol used. These constraints are mostly inherited by the permissionless setting, e.g. computational power in proof-of-work, and become serious obstacles in a resource-constrained IoT environment. Moreover, consensus protocols with low throughput or high latency are not suitable for IoT networks where massive volumes of data are generated. Thus, in this paper we focus on permissioned blockchain platforms and investigate the consensus protocols used, aiming at evaluating their performance and fault tolerance as the main selection criteria for (in principle highly insecure) IoT ecosystem. The results of the paper provide new insights on the essential differences of various consensus protocols and their capacity to meet IoT needs.

**Index Terms**—Permissioned blockchains, consensus protocols, cyber-security, fault tolerance, Internet of things.

## I. INTRODUCTION

The blockchain is an innovative technology that has proved to bring considerable advantages in many areas by inherently providing security, interoperability and sustainability; it can be used as an independent root-of-trust in an distributed (possibly adversarial) setting to allow a set of entities mutually trust each other. For this reason, the blockchain is considered as a decentralized technology for the sharing of information and performing transactions in a secure manner. In contrast to permissionless blockchains, in which any node can participate in the network (Ethereum, Bitcoin, etc.), the permissioned blockchains are characterized by the fact that the participating nodes who regulate and amend the state of the ledger can be identified and can be held accountable for their actions. The key idea behind this concept is to address the privacy needs and other requirements of distributed applications, as well as to provide convenient access control mechanisms with some means of identification.

The current state-of-the-art in permissioned blockchains is rather to adopt the design decisions made in permissionless architectures [1]. Unfortunately, in most cases this approach leads to sub-optimal performance since a well-defined protocol under the rules of a public ledger, often with financial orientation, might not be well addressing the challenges of private networks and IoT environments. The transactions' fees, the (high) processing time and (low) networks' throughput should not be inherited by the consensus protocols of permissioned blockchains. Apart from the above constraints, the IoT devices should not be part of the consensus process themselves, due to the fact that several consensus protocols demand considerable computational resources that smart devices do not possess. These issues have been identified in numerous permissioned blockchains [1], where the design decisions inherited by the permissionless settings pose great threats to the security, performance, and eventually the blockchain solution's suitability in demanding IoT applications.

Apart from several attacks in blockchains and distributed ledgers, like double-spending and selfish mining, permissioned blockchains often rely on low-level trust models for validation that are originating from the consensus protocol and are hard to be adapted to the operation of smart contracts [2]. In addition, nodes involved in the consensus protocol might stop working, start behaving in a malicious manner, or acting selfishly against common goals and thus undermine the security of the protocol. Therefore, the consensus nodes should execute a fault-tolerant protocol to safeguard the integrity of the transactions as well as the total order in which they are included in the blockchain in order to deliver a continuous service.

Motivated by the above concerns, and the need for enforcing accountability in many applications, this paper focuses on permissioned blockchains and analyses the properties of a large number of state-of-the-art consensus protocols. Accountability is achieved either directly by immediately proving that a node behaved in a dishonest way, or indirectly by inferring possible malicious behavior using side information at some later stage; trust management schemes can be used to expose long-term malicious trends that are difficult to distinguish at any given time if considered independently. Our work complements the reports by Cachin [3], Xiao [4], Zheng [5], Wu [6], Salimitari [7], Lao [8], and Ferdous [9], by considerably extending the number of blockchain platforms and consensus protocols that are studied, by providing quantitative information about their



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 786698. The work reflects only the authors' view and the Agency is not responsible for any use that may be made of the information it contains.

performance (i.e. throughput and latency) instead of qualitative information that is given in all works (thus, allowing to easily consider issues related to scalability), and by evaluating their suitability for IoT applications also taking into account their resilience to adversarial faults. The comparative evaluation of the various consensus protocols suggests that only few of them meet IoT needs, in addition to achieving high security.

The paper is next organized as follows. Section II introduces the basics about consensus protocols’ security against faults. In Section III and IV we give a detailed analysis of the most prominent crash fault tolerant and Byzantine fault tolerant consensus protocols for permissioned blockchain platforms. The comparative evaluation of the various protocols is provided in Section IV-M, whereas Section VI concludes the paper and discusses future research steps.

## II. FAULT TOLERANCE PROPERTIES

Distributed systems are often disciplined by a set of clients and services, where each service utilizes one or more servers to extract information or execute operations that are requested from the clients. Using a central server is the easiest way to fulfill the necessary needs and the implementation of a service, but it poses a major concern about the security of an IoT network, due to the fact that it becomes a *single point of failure* (SPoF). Thus, to avoid centralized faults, multiple servers should be deployed to implement a Fault Tolerant service with *state machine replication* (SMR) [10].

The pioneering work of Lamport [11], who introduced the Byzantine Agreement, triggered the research of developing algorithms in order to exploit and construct resilient distributed systems. Nonetheless, several blockchain systems deviate from the classical SMR in crucial ways. Many distributed applications run simultaneously and can be deployed at any time, even if the embedded application code is un-trusted or occasionally malicious. The key idea to provide security in blockchains is to reach agreement on a single request from a client, which is the core functionality of a consensus protocol. In the context of blockchains, a consensus protocol or as it is commonly known as an “atomic broadcast”, provides a total order of the disseminated messages and propagates them to the network peers.

### A. Crash fault tolerance

An atomic broadcast certifies that all the legitimate nodes output or deliver the identical array of messages by means of the deliver event. Accurately, considering a set of  $n$  nodes in the network, it certifies that the properties validity, integrity and total order are fulfilled [12].

The way to achieve consensus (i.e. to realize atomic broadcast) in distributed systems that are vulnerable to  $t < n/2$  node crashes is to adopt consensus protocols known as the *viewstamped replication* (VSR) [13] and *Paxos* [14] family of protocols. To provide security, this family is characterized by the same rules. In each round a leader is elected or voted to create a new block and if the ongoing leader crashes or even if the nodes in the network suspect that the leader has crashed,

the leader is reinstated by proceeding to the next round. This family of protocols is known today as *crash fault tolerant* (CFT) consensus protocols and it guarantees that a set of failing nodes  $t < n/2$  does not impact the system.

### B. Byzantine fault tolerance

Consensus protocols with the purpose of tolerating byzantine nodes, which are subverted by a malicious actor and avert the common goal of reaching agreement, have recently emerged. In the *Byzantine fault tolerant* (BFT) consensus protocols family, the most common protocol is the *practical Byzantine fault tolerant* (PBFT) [15], which can be displayed as a blossom of the *VSR/Paxos* [13], [14]. In a network comprised of a set of  $n$  nodes, the PBFT consensus protocol [15] can tolerate  $f < n/3$  subverted nodes using a progression of rounds with a unique leader within each round. Under the assumption that the BFT protocols seem to be more secure than CFT, vast research work has focused on the improvement of the PBFT with *BFT-SMaRT* [16] to be considered one of the most advanced and scalable BFT consensus protocols.

## III. CRASH FAULT TOLERANT CONSENSUS

Fault-Tolerant algorithms have acquired momentous observation over the years [17]. Some of them provide significant results of solving consensus, in an optimum way, by identifying bounds on the security beneath different models. In this section, we survey recent CFT consensus protocols, focusing on results for permissioned, IoT-based blockchain platforms. However, these protocols cannot endure malicious activities but they can only tolerate 50% of the network crashes, making their adoption for IoT networks not a very appropriate choice.

### A. Kafka / Zookeeper: Hyperledger Fabric, Corda

Kafka [18] is a distributed publish-subscribe streaming platform, embraced by *Fabric* [2] and experimentally by *Corda* [19] as a cluster of ordering/notary service. At high level, the conceptual configuration of *Kafka* is identical with the leader – follower setting. The transactions, which *Kafka* calls *messages*, are replicated from the leader to its followers and if the leader crashes, then one of the followers takes command. This action ensures the crash-fault-tolerance of the network. The administration of *Kafka* [18], including the systematization of tasks, the election process, the cluster association, among others, is organized by *Zookeeper* [20]. Despite the fact that *Kafka* is CFT, *Zookeeper* by drifting *Kafka* becomes somehow centralized and thus the protocol is not advised to be deployed in a large network. In simple words, the ordering service of *Fabric* is executed by distinct organizations and *Kafka* demands one of them to run the cluster. This means that all the ordering nodes are attracted to the same cluster, which is under the jurisdiction of a single organization. Thus, this concept does not assure much in terms of decentralization, even though that the protocol achieves a significant throughput with very low latency [2]. In the case of *Corda*, the enterprise described a prototype of a *Kafka*-based CFT notary service [19]. Unfortunately,

the experimental implementation of `Kafka` in `Corda` is not expected to be deployed soon.

### B. Raft: Hyperledger Fabric, Sawtooth, Quorum, Corda

`Raft` [21] is a CFT consensus protocol, alternative to `Paxos` [14], implemented in `Hyperledger Fabric` [2], `Hyperledger Sawtooth` [22], `Quorum` [23] and `Corda` [24]. As a voting system, `Raft` realizes the leader-follower model, in which each node is characterized as *leader*, *follower* or *candidate*. The consensus process is parted into the following sub-problems: a) the leader election, with the initiation of the protocol or in case the ongoing leader crashes; b) the log replication, in which the leader affirms log entries, replicates them and compels all the followers to acknowledge them and c) the safety of the protocol, concerning the regulations imposed on the leader-election process to ensure this security property [21]. The time in `Raft` is separated into *terms*. The terms are measured with increasing integers and initiated with an election. If a candidate wins the election, then it takes the leader's place for this term. The leader disseminates systematic heartbeat messages to its followers and if, for a preordained time period a candidate obtains no heartbeat messages, then it is presumed that the leader has crashed, which means that a new election process has to be initiated.

Although `Raft` is a fast CFT consensus protocol, it provides low and medium performance if it is implemented in `Corda` [25] and `Quorum` [26] (respectively). On the other hand, in the new version of `Fabric` (`Fabric v2.0`) [27], in which `Raft` is the recommended protocol, it is possible to achieve twice the transaction throughput of `Kafka`'s and thousands of transactions in real world scenarios [28], [27] which makes `Fabric` with `Raft` a very good candidate for IoT applications.

## IV. BYZANTINE FAULT TOLERANT CONSENSUS

A primary case of failures that is essential to blockchains, is the malicious behavior caused by an adversary. The failure of nodes is stated, not only on unintentional crashes, but also on contemplated attacks in the system, that attempt to defeat the security enhancements by causing at least forks in the ledger. In the context of blockchains, the consensus mechanisms are required to address malicious, selfish and generally any set of nodes that attempt to alter the protocol, by ensuring the safety and the global state coherence. For this reason, an adoption of a BFT protocol with high performance metrics is more suitable for an IoT environment.

### A. PBFT: Hyperledger Sawtooth

`PBFT` is the first BFT protocol based on SMR and was developed by Castro and Liskov [15] in 1999. The protocol has become a synonym of BFT and has attained immense concern in distributed consensus. Each participating node is a *validating replica* (VR); one of them is named to be the leader. The *leader validating replica* (LVR) receives a transaction request from a client, who demands an execution of an operation. Then, the LVR validates the transaction and

disseminates this request to the other VRs. In a short period of time, called "batch-timeout", or after a plethora of ordered pending transactions, called "batch-size", a block is created by the LVR. Then, the block is disseminated to the other VRs to reach consensus, initializing at the same time the **pre-prepare**, **prepare** and **commit** phase of consensus. If  $2f+1$  VRs come to an agreement upon a decision, then each VR appends this block as the next block in its private ledger. Although the protocol provides the desirable performance of over 78000 tps as shown in [16], it lacks scalability supporting only few tens of participants in the network; demanding high communication complexity due to the high number of messages that have to be transferred [29], [30]. These reasons suggest that the protocol should be deployed in a small IoT environment, for example in a *small office/home office* (SOHO) environment and not in a large network where the high number of nodes can act as bottleneck.

### B. IBFT: Quorum, Hyperledger Besu, Autonomy

*Istanbul Byzantine fault tolerant* (IBFT) [31] is a quite appealing *proof-of-authority* (PoA) mechanism, which is emanated from the `PBFT` and inherits its security properties from it, by employing the three-phase consensus. In a similar way to `PBFT`, the VRs elect a LVR to create a proposed block and disseminate it to the network, along with a **pre-prepare** message. The VRs, upon accepting this **pre-prepare** message and to be certain that they are on the same sequence, they enter the **pre-prepared** phase and disseminate a **prepare** message. The LVR, in its turn, while gathering a number of  $2f+1$  **prepare** messages from the VRs, enters the **prepared** phase and disseminates a **commit** message. This means that the LVR accepts the proposed block, while declaring that this block is going to be inserted to its ledger. Then, the VRs, upon receiving  $2f+1$  **commit** messages, they enter the **committed** phase and accept the proposed block. Although, IBFT is similar to `PBFT`, the later needs some tweaks to be used in blockchains. In IBFT, the fact that there is no particular client to demand results by sending requests, makes all the VRs to be recognized as such. The performance evaluation of Baliga et. al., [26] for using IBFT and `Raft` in the `Quorum` platform is not peer-reviewed, but it shows that when private smart contracts are used, in which it is required further encryption - decryption functions and further communication overhead among peers, it can be achieved a throughput of 600 - 650 tps with a latency of around 4.5sec.

### C. BFT-SMaRT: Hyperledger Fabric, Corda, Symbiont

Bessani's protocol is a java-based library implemented in various platforms such as `Hyperledger Fabric` [2], `R3 Corda` [19] (experimentally) and `Symbiont Assembly` [32]. In the absence of malicious VRs, `BFT-SMaRT` [16] achieves consensus using the message model of `PBFT` [15], but in their presence the protocol forces the network to elect a new LVR and execute the messaging pattern, which is described in [16]. The differences between the two BFT protocols is firstly that `BFT-SMaRT` has enhanced reliability and

multi-core processing concerning the appraisal of signatures. BFT-SMaRT realizes a modular consensus protocol, which is not embodied inside of the SMR while the state transfer and reconfiguration modules are clearly detached from the agreement method [16]. The supported reconfiguration makes the protocol to differ from previous BFT systems, in which the size of the network could not expand or reduced, and allows VRs to join or leave the network on-the-fly. In the context of *Fabric*, the consensus protocol also implements Wheat [33], a vote assignment model (as the VRs recognize each other), to achieve the desirable for IoT networks performance without imperiling the safety and the stability of the consensus. Although, the protocol is not in a stable mode if it is implemented in *Corda*, in *Fabric*'s and *Symbiont*'s case the desirable performance for IoT deployment has been achieved far more than necessary.

#### D. RBFT: *Ontology*, *Hyperledger Indy*

RBFT which stands for *redundant Byzantine fault tolerance* is a consensus algorithm proposed by Aublin *et al.* [34] and implemented in *Ontology* [35] and in *Indy* [36]. The recent BFT protocols elect in each instance a primary replica, herein called LVR, which demonstrates how all the incoming requests will be ordered. RBFT introduces a new approach, running multiple instances concurrently, in order to identify and quickly replace malicious LVRs, with the reduction of throughput's degradation caused by their presence as its aftereffect. While multiple instances are executed and order requests, their performance is strictly monitored and only the requests that have been ordered by the master instance are executed. The detection of a malicious master LVR, which wants to smartly and willingly degrade the performance of the protocol, results from comparing its performance with all the other LVRs' performances. This comparison takes place due to the fact that in BFT protocols it is hard for VRs to presume the throughput that a non-malicious LVR would have. In the case, where a LVR is faster than the master LVR, then a new LVR is elected in each instance and the master LVR is marked as malicious. In RBFT [34], when malicious faults occur, the throughput is degraded only by 3%, while the degradation of other existing BFT protocols, under the same faults, is at the best of circumstances equivalent to 78%. This property makes the protocol a quite appealing mechanism to be used in IoT environments; if it is combined with a platform with high performance.

#### E. VBFT: *Ontology*

VBFT [37], the core consensus protocol of *Ontology*, [35], relies on *verifiable random functions* (VRF) [38] to introduce randomness and is combined with *proof-of-stake* (PoS) and Byzantine tolerance to provide resistance against malicious acts. The whole network consists of two different types: The consensus nodes, where the stake of each participant has significant impact; and the consensus candidates, in which the nodes do not aid the consensus process, but validate consensus blocks and update their ledger. Due to the

randomness provided by the VRF function, the VBFT protocol selects different sets of nodes, which are difficult to predict and each one of them is assigned a different work to perform. At first, each proposal node creates and proposes a new block. The proposed blocks are collected from the verification nodes, which verify and independently vote for them depending on their highest priority. Then, the confirmation nodes confirm the results that were provided by the verification nodes and finally establish a consensus result. Upon the reception of an established result, a new round begins. The protocol is only implemented in the *Ontology* [35] platform – a public, high performance and scalable permissioned blockchain - that enables smart contracts for different business requirements.

#### F. Tendermint: *Hyperledger Burrow*, *Autonity*, *Ethermint*

Tendermint [39], which is majorly adopted in *Burrow* [40], *Autonity* [41] and *Ethermint* [42], is designed as a deterministic protocol under a partially synchronous communication. Although the Tendermint consensus has similarities to other BFT protocols, the voting power in this protocol differs and is defined based on each node's stake. Particularly in the propose phase, a VR is deterministically selected, according to the rate of its stake, to propose a block for a specific height. Then, for a block to be committed, the pre-vote, pre-commit and the commit phases have to follow in a way similar to most BFT protocols.

Furthermore, a concept called *locks* or *polka* is included in Tendermint's terminology. More precisely, in the pre-commit phase, if a proposed block gathers more than the  $2/3$  of the VRs' pre-votes for a block, then a lock on the proposed VR occurs. For various reasons, a new block may fail to be committed in a specific height. In this case, the protocol moves on to a new round with a new VR to propose a new block for that height. In such a round, the new proposer may be locked in a block from a preceding round. Then the proposed-locked block is the same as before and in its proposal a *proof-of-lock* (PoL) is added. The PoL [39] is a collection of pre-votes from the VRs, concerning the situation of a proposed block, whether that is locked or unlocked.

In the pre-vote phase, each locked VR signs and disseminates a pre-vote concerning its locked block. Frequently, due to network asynchrony, the VRs may not obtain any proposal or even worse, obtain an invalid block. In such a case, they sign and then disseminate a nil pre-vote. Generally, a nil vote, depending on the phase that is taking place, is a vote to either move to the next round or to unlock a block and a situation in a later round, when another block is locked for the same height. In the pre-commit phase, if more than  $2/3$  from the pre-votes are nil votes, then the VR and the block are considered unlocked and a propose phase of a new round is initiated. Otherwise, the VRs enter the commit phase. Although Tendermint is a secure BFT consensus protocol, its implementation in *Ethermint* and *Burrow* provides low and medium transaction throughput (respectively) making the protocol not the best option for IoT environments.

### G. Exonum consensus: Exonum

The Exonum consensus [43], a customized BFT algorithm, is built by Bitfury and implemented only to the Exonum platform. The network comprises two different kinds of nodes, with each kind to have been assigned a different work to perform; on the one hand the Auditors, which do not assist in reaching consensus but can read the blocks' transactions and on the other hand the validators, herein called VRs, which participate in the consensus process. As it commonly happens in a BFT protocol, each round is initiated with an elected LVR who sends a proposal to be embodied in the blockchain. The three-phase consensus (pre-vote, pre-commit and commit) is almost the same for the Exonum consensus. Similarly to Tendermint, a special concept called "lock" is added to the protocol. The lock defines that the VR, whose pre-vote block has gather more than  $2/3$  of the approvals from the network, gives up voting for other VRs' proposals and locks on its own, without giving any permission to change it. Then, the VR broadcasts a pre-vote message expressing its locked condition on the specific proposal. With high throughput, low latency and the support of smart contracts, Exonum makes a good candidate to be implemented in IoT networks if the concept of "locks" is connected with IoT-based criteria.

### H. DPoS: EOSIO, Bitshares, Tron, Tezos, Lisk

*Delegated proof-of-stake* (DPoS), a reputation-based consensus protocol, is implemented in Bitshares [44], EOSIO [45], Tron [46], Tezos [47], and Lisk [48]. In DPoS, the nodes vote with reputation scores to choose a class of delegates that will be assigned to create blocks. In each round, among the set of delegates, a leader is selected in a way defined by the respective distributed ledger. The leader is incentivized to follow the protocol and get rewarded upon the creation of an honest block and penalized or blacklisted in any other case. Among the delegates, a contest on which one of them is going to be included in the validation set takes place, with each delegate promising various levels of rewards to its voters if it is elected as a leader. Each ledger follows its own parameters and by using a small number of validators, finality can be achieved promptly. The performance evaluation of DPoS varies across different platforms, with EOSIO and Bitshares to provide a desirable, suitable and high performance network for IoT environments if the stake is replaced with IoT-based criteria.

### I. Clique / Aura: Quorum, Parity, Autonity

Clique and Aura are POA protocols [29] developed to supposedly displace the computational consumption in permissioned blockchains and to increase the transaction throughput. Clique is implemented in Quorum [23] and Autonity [41]. The time in Clique proceeds in epochs, where during each epoch the current mining leader and multiple other authorities are allowed to propose a block after a determined number of epochs. Aura, on the other hand, is implemented in Parity [49] and does not commit the proposed block at once, but rather adds an extra round, called "block acceptance",

to examine if the received block is the same to all the authorities. In cases where the leader has not proposed any block (even empty) or has proposed multiple and contrasting blocks to various authorities, then the leader's position is under vote. If the leader is emerged as malicious then it is removed. An academic research [29] showed that under the presence of byzantine authorities Clique and Aura may fork and thus result in the violation of the protocol's safety and reduction of the network's throughput. Clique and Aura provide high Byzantine fault tolerance of  $2f+1$  nodes, but as consensus protocols might encounter network delays and block confirmation time over 5sec, which may not be admissible in delay-delicate IoT environments.

### J. DBFT: NEO, Ontology

Developed from the NEO team [50] and later adjusted to the Ontology platform [35], the *delegated Byzantine fault tolerance* (DBFT) protocol follows the structure of PBFT and DPoS, providing thus the best of both worlds. From PBFT, the protocol inherits the three-phase consensus without the need for all the nodes to participate in it. With this adjustment the protocol becomes stable and averts malicious forks. On the other hand, the voting system, in which only the delegates can execute the consensus process, is inherited from the DPoS protocol, providing thus a sampling of the most trusted entities in the network. The derived advantages of the combination of the two are: the instant transaction finality without forks and a fast voting-based method to elect the delegates. The high transaction throughput and the security that the protocol provides are enticing features for IoT applications. However, in Ontology and in NEO platform, the block confirmation time is 5–20sec, which is beyond the appropriate restraints and makes the protocol not so applicable for IoT devices.

### K. YAC: Hyperledger Iroha

*Yet another consensus* (YAC) [51] is a novel protocol implemented in Hyperledger Iroha [52] with its priority the mobile app development using the simplest possible construction. In the YAC protocol, the validation process is not part of the consensus, but it is rather based on the transaction flow of Iroha. In order to be considered valid, each transaction has to pass at first the stateless validation —in which the transaction's signature and format are checked— and afterwards the stateful validation that takes place in a slower form, after the ordering process. The procedures performed by the protocol are limited in ordering and consensus; The first is executed by the ordering service and the later by an ordered list of peers, in which the first peer is considered to be the leader. Although, the Iroha's transaction flow to an extent reminds us of Fabric's, here only the verified transactions are included in a block. Upon the reception of the ordered sequence of transactions in a pattern of proposals, the peers perform the stateful validation check, remove invalid transactions from the sequence of transaction, create a proposed block, vote the proposed block through signing and then disseminate it to the leader. After receiving  $2/3$  of the proposed-voted

blocks, the leader sends to all the consensus peers a commit message manifesting which block should be accepted to their chains. Although Iroha’s throughput is satisfiable for IoT environments, the block confirmation time is slightly over from other favorable permissioned blockchains due to the restraints that small vote delays enforce on the peers in order to come to an agreement on a specific proposal.

#### L. PoET: Sawtooth

To solve the problem of Byzantine agreement and to avoid the wasteful computational power that Nakamoto’s consensus utilizes [53], *proof-of-elapsed time* (PoET) implements a lottery-based algorithm to achieve fairness, investment and verification in the leader election process. The peers in Sawtooth [22] are thoughtfully elected to execute requests after waiting an indefinite period of time [54]. The peer with the lowest waiting time creates the new block. To certify that each peer’s time has actually elapsed, the protocol demands the whole critical process to be performed in a private memory area called *trusted executed environment* (TEE), i.e. Intel SGX [55]. Specifically, the TEE’s utility, among others, is to provide a cohesion proof that results from a trusted function, often called “enclave”, by means of remote attestation in order to establish trust on the consensus network. The protocol is BFT but malicious adversaries can manipulate the network if they compromise participating nodes beyond the threshold of  $\Theta(\frac{\log \log n}{\log n})$  [56]. The protocol possesses the required security and performance criteria to be deployed in IoT environments, but without integration of the SGX hardware, PoET’s security degrades into the family of CFT protocols.

#### M. RPCA: Ripple

The *Ripple consensus algorithm* (RPCA), was introduced in 2014 by Schwartz *et al.* [57], as the fundamental protocol for a secure and fast real-time cryptocurrency-based system to transfer remittances without the support of smart contracts. In the RPCA, the nodes propose the assembled transactions from their clients to the validators to reach consensus. The time in Ripple is measured in *epochs*, with each epoch involving several rounds of transactions’ refinement. Initially, with the use of a *unique node list* (UNL), each node identifies all the nodes with which it can instantaneously interact, exchange messages and trust. In a vast network, such as Ripple, trust does not actually mean that each node in a UNL is trusted, but rather that this node will not attempt to circumvent the network with votes on inaccurate transactions. A transaction with a minimum threshold of yes-votes is forwarded to the next round, while others, that do not exceed the threshold, are either deserted; if malicious, or embodied in a *candidate set* of new and not yet applied transactions on the ledger. The transactions that in the final round exceed with positive votes with a percentage of 80% of the node’s UNL, are included to the ledger, while the candidate set is forwarded for consensus on the next ledger. Taking into account that Ripple is a semi-permissioned system, its cryptocurrency-based nature, its overall performance and the fact that the percentage of the

TABLE I  
CRITERIA FOR THE SUITABILITY OF CONSENSUS PROTOCOLS IN IoT

Criteria	Suitability	Symbol
FT and LP	0 %	○
CFT and MP	25 %	◐
BFT and MP	50 %	◑
CFT and HP	75 %	◒
BFT and HP	100 %	◓

tolerated malicious validators in each node’s UNL is 20%; the protocol is not a good option for IoT deployment.

### V. COMPARATIVE EVALUATION

The consensus protocols presented in the previous sections have been incorporated into various permissioned blockchain platforms. The transaction throughput and latency are the most critical factors to consider in IoT deployments, as it is required a transaction to be broadcast and appended into a block in a matter of seconds or even less. According to [7], throughput and latency are mapped to a three-valued scale, namely *low*, *medium*, and *high*, where the thresholds that define throughput are 100 and 1000 *transactions per second* (TPS) respectively; in addition, latency characterization relies on the time units (milliseconds, seconds, or minutes). In this paper, we follow the same three-valued scale but the throughput (resp. latency) thresholds have been adjusted to 500 and 1500 TPS (resp. 1sec and 10sec) respectively to meet IoT applications’ requirements.

Table I, explains the symbols used below for characterizing the consensus’ suitability in IoT networks. With ◓ the most prominent BFT protocols having *high performance* (HP) are denoted, whereas ◑ is used for BFT protocols with *medium performance* (MP). On the contrary, high performance CFT protocols are appropriate for IoT applications only under the assumption of non-adversarial and honest environment. These assumptions are quite restrictive, focusing on resiliency only, and hence these protocols’ suitability ◑ is not the maximum possible; for the same reasons, the suitability of CFT protocols with a medium performance is ◐. Finally, fault tolerant and *low performance* (LP) protocols are denoted with ○, implying that they are not suitable to IoT environments.

Next, a comparative evaluation of 13 consensus protocols and their implementation into blockchain platforms is shown in Table II; the performance metrics used to determine IoT suitability are fault tolerance, the adoption under permissioned blockchain platforms, transactions’ performance, and latency. The transaction’s performance is measured in TPS, while the values reported for latency concern either the block confirmation time or the time needed for a transaction to be ordered or notarized, e.g. in Fabric, Corda, etc. The analysis of centralized and PoW based consensus protocols are not included in our survey and in Table II, as they are either designed for just facilitating development or they are computationally expensive. The approach used to compute the

TABLE II  
COMPARATIVE EVALUATION OF BLOCKCHAIN PLATFORMS' CONSENSUS PROTOCOLS AGAINST THE NEEDS OF IOT

Consensus Protocol	Fault Tolerance	Blockchain Platforms	Throughput (TPS)	Latency – Block Confirmation Time	Evaluation Source	Overall Performance	IoT Applicability
Aura	BFT: $2f + 1$	Parity	35 45	3 7sec	[58]	Low	○
BFT-SMaRt	BFT: $3f + 1$	Fabric	> 10K	0.5sec	[33]	High	●
		Symbiont	80K	< 1sec	[3]	High	●
DBFT	BFT: $3f + 1$	NEO	< 1K	15 20sec	[50]	Low	○
		EOSIO	1K 6K	< 1sec	[45]	High	●
DPoS	BFT: $3f + 1$	Bitshares	100K	1sec	[44]	High	●
		Tron	> 2K	3sec	[46]	Med	◐
		Lisk	2.5	6min	[59]	Low	○
		Tezos	≈ 40	≈ 30min	[47]	Low	○
Exonum	BFT: $3f + 1$	Exonum	≈ 5K	0.5sec	[60]	High	●
IBFT	BFT: $3f + 1$	Quorum	≈ 600	4.5sec	[26]	Med	◐
		Fabric	3.5K	< 1sec	[2]	High	●
Kafka	CFT: $2t + 1$	Corda	3K 15K	1sec	[19], [61]	High	●
PoET-SGX	BFT: $\Theta(\frac{\log \log n}{\log n})$	Sawtooth	1K 2.3K	< 1sec	[25]	High	●
		Corda	100 200	1sec	[25]	Low	○
Raft	CFT: $2t + 1$	Quorum	≈ 650	4.5sec	[26]	Med	◐
		Fabric	≈ 7K	< 1sec	[28]	High	●
Ripple	BFT: $4f + 1$	Ripple	1.5K	4sec	[62]	Med	◐
Tendermint	BFT: $3f + 1$	Ethermint	200 800	< 1sec	[42]	Med	◐
		Burrow	> 400	n/a	[63]	Low	○
VBFT	BFT: $3f + 1$	Ontology	> 3K	5 10sec	[35]	Med	◐
YAC	BFT: $3f + 1$	Iroha	several 1K	< 3sec	[51], [64]	Med	◐

overall performance in Table II is that it equals the minimum of the scores achieved by a protocol's throughput and latency.

Although, the above features are evaluated under different setups for different blockchain platforms, the performance of several consensus protocols is not mentioned in their white papers. It is expected that, if implemented in a secure, fast, and scalable permissioned blockchain, a high performance consensus protocol should have a promising throughput and low latency. For example, *Sawtooth* is considered to be a high performance blockchain platform. However, it is not thoroughly tested under different setups in a large scale environment yet. For this reason, such protocols are not currently included in Table II. The main outcome of the analysis is that even by replacing the monetary concepts with IoT-based criteria, a small number of the proposed consensus protocols meet all the performance and security requirements. With the exception of *Sawtooth* that is hardware supported by a TEE in its *PoET-SGX* consensus, the rest of the BFT protocols (BFT-SMaRt, DPoS, and Exonum) achieve quite a high throughput that can reach 100K TPS and latency as low as 0.5sec.

## VI. CONCLUSIONS

In this paper, the potential of using fault-tolerant consensus protocols in IoT networks is investigated. The protocols are analyzed in terms of their performance under different permissioned blockchain platforms and their tolerance against faults, including adversarial ones — a property that is highly desirable

in the generally vulnerable IoT ecosystem. The integrated BFT protocols in *Hyperledger Fabric*, seem to offer a good option for IoT networks when it comes to permissioned blockchains.

Remarkably, several of the surveyed consensus protocols are under development and not thoroughly tested at the time of writing; for this reason, they have not been included in the comparison. Apart from the performance and fault tolerance that are explored in this paper, several other limitations are faced when blockchain technology is applied in IoT networks. Limitations such as scalability, data privacy, and several other aspects that are part of our going research and will be included in a forthcoming work.

## REFERENCES

- [1] M. Vukolic, "Rethinking permissioned blockchains," in *Proc. of ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, 2017, p. 3–7.
- [2] E. Androulaki *et al.*, "Hyperledger fabric: A distributed operating system for permissioned blockchains," in *Proc. of 13th EuroSys Conf.*, ser. EuroSys '18, 2018, pp. 30:1–30:15.
- [3] C. Cachin and M. Vukolic, "Blockchain Consensus Protocols in the Wild (Keynote Talk)," in *31st Int'l Symposium on Distributed Computing (DISC 2017)*, vol. 91, 2017, pp. 1:1–1:16.
- [4] Y. Xiao *et al.*, "A survey of distributed consensus protocols for blockchain networks," *IEEE Communications Surveys Tutorials*, pp. 1–1, 2020.
- [5] Z. Zheng *et al.*, "An overview of blockchain technology: Architecture, consensus, and future trends," in *2017 IEEE Int'l Congress on Big Data (BigData Congress)*, June 2017, pp. 557–564.
- [6] M. Wu *et al.*, "A comprehensive survey of blockchain: From theory to iot applications and beyond," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8114–8154, Oct 2019.

- [7] M. Salimitari and M. Chatterjee, "An overview of blockchain and consensus protocols for iot networks," *CoRR*, vol. abs/1809.05613, 2018.
- [8] L. Lao *et al.*, "A survey of iot applications in blockchain systems: Architecture, consensus, and traffic modeling," in *CSUR*, 2020.
- [9] M. Sadek Ferdous *et al.*, "Blockchain Consensus Algorithms: A Survey," *arXiv e-prints*, p. arXiv:2001.07091, Jan 2020.
- [10] F. B. Schneider, "Implementing fault-tolerant services using the state machine approach: A tutorial," *ACM Comput. Surv.*, vol. 22, no. 4, pp. 299–319, 1990.
- [11] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, Jul. 1982.
- [12] V. Hadzilacos and S. Toueg, "Fault-tolerant broadcasts and related problems," in *Distributed Systems (2Nd Ed.)*, S. Mullender, Ed., 1993, pp. 97–145.
- [13] B. M. Oki and B. H. Liskov, "Viewstamped replication: A new primary copy method to support highly-available distributed systems," in *Proc. of 7th Annual ACM Symposium on Principles of Distributed Computing*, 1988, pp. 8–17.
- [14] L. Lamport, "The part-time parliament," *ACM Trans. Comput. Syst.*, vol. 16, no. 2, pp. 133–169, May 1998.
- [15] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, 2002.
- [16] A. Bessani, J. Sousa, and E. E. P. Alchieri, "State machine replication for the masses with bft-smart," in *2014 44th Annual IEEE/IFIP Int'l Conf. on Dependable Systems and Networks*, 2014, pp. 355–362.
- [17] N. A. Lynch, *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [18] "Kafka 2.3 documentation," 2001. [Online]. Available: <https://kafka.apache.org/documentation/>
- [19] "Corda documentation." [Online]. Available: <https://docs.corda.net/design/kafka-notary/design.html>
- [20] P. Hunt *et al.*, "Zookeeper: Wait-free coordination for internet-scale systems," in *USENIX ATC'10: Proc. of 2010 USENIX Annual Technical Conf.*, Jun 23-25 2010, pp. 11–11.
- [21] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in *2014 USENIX Annual Technical Conf. (USENIX ATC 14)*, Jun. 2014, pp. 305–319.
- [22] "Sawtooth v1.2.4 documentation," 2018. [Online]. Available: <https://sawtooth.hyperledger.org/docs/core/releases/latest/contents.html>
- [23] "Quorum documentation," 2018. [Online]. Available: <http://docs.goquorum.com/en/latest/>
- [24] "R3 primer series 1." [Online]. Available: <https://www.r3.com/wp-content/uploads/2019/01/R3-Quick-Facts.pdf>
- [25] B. Ampel, M. Patton, and H. Chen, "Performance modeling of hyperledger sawtooth blockchain," in *2019 IEEE Int'l Conf. on Intelligence and Security Informatics (ISI)*, July 2019, pp. 59–61.
- [26] A. Baliga *et al.*, "Performance evaluation of the quorum blockchain platform," *CoRR*, vol. abs/1809.03421, 2018.
- [27] "Hyperledger fabric documentation." [Online]. Available: <https://hyperledger-fabric.readthedocs.io/en/release-2.0/whatsnew.html>
- [28] C. Ferris, "Does hyperledger fabric perform at scale?" [Online]. Available: <https://www.ibm.com/blogs/blockchain/2019/04/does-hyperledger-fabric-perform-at-scale/>
- [29] S. D. Angelis *et al.*, "Pbft vs proof-of-authority: applying the cap theorem to permissioned blockchain," in *Italian Conf. on Cyber Security (06/02/18)*, January 2018.
- [30] J. Liu *et al.*, "Scalable byzantine consensus via hardware-assisted secret sharing," *IEEE Transactions on Computers*, vol. 68, no. 1, pp. 139–151, 2019.
- [31] "Istanbul byzantine fault tolerance," 2017. [Online]. Available: <https://github.com/ethereum/EIPs/issues/650>
- [32] "Symbiont documentation," 2019. [Online]. Available: <https://symbiont.io/>
- [33] J. Sousa, A. Bessani, and M. Vukolic, "A byzantine fault-tolerant ordering service for the hyperledger fabric blockchain platform," in [35]
- [34] P. Aublin, S. B. Mokhtar, and V. Quéma, "Rbft: Redundant byzantine fault tolerance," in *2013 IEEE 33rd Int'l Conf. on Distributed Computing Systems*, July 2013, pp. 297–306.
- [35] "Indy documentation." [Online]. Available: <https://indy.readthedocs.io/en/latest/index.html>
- [36] Honglei-Cong, "Vbft introduction," 2018. [Online]. Available: <https://github.com/ontio/documentation/blob/master/vbft-intro/vbft-intro.md>
- [37] S. Micali, M. Rabin, and S. Vadhan, "Verifiable random functions," in *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, Oct 1999, pp. 120–130.
- [38] J. Kwon, "Tendermint: Consensus without mining," *Draft v. 0.6, fall*, vol. 1, no. 11, 2014.
- [39] "Hyperledger burrow documentation." [Online]. Available: <https://github.com/hyperledger/burrow/blob/master/README.md>
- [40] "Autonity documentation." [Online]. Available: <https://docs.autonity.io/deploy/index.html>
- [41] "A beginner's guide to ethermint." [Online]. Available: <https://blog.cosmos.network/a-beginners-guide-to-ethermint-38ee15f8a6f4>
- [42] Y. Yanovich *et al.*, "Exonum: Byzantine fault tolerant protocol for blockchains," *bitfury.com*, pp. 1–36, 2018.
- [43] "Bitshares documentation." [Online]. Available: <https://bitshares.org/technology/industrial-performance-scalability/>
- [44] "Eosio documentation." [Online]. Available: <https://eos.io/news/askblockone-what-are-the-benefits-of-the-eosio-infrastructure/>
- [45] "Tron documentation." [Online]. Available: <https://tron.network/>
- [46] "Tezos (xtz)." [Online]. Available: <https://blockfury.com/tezos-xtz/>
- [47] "Lisk documentation." [Online]. Available: <https://lisk.io/documentation/lisk-sdk/index.html>
- [48] "Parity tech documentation," 2019. [Online]. Available: <https://wiki.parity.io/>
- [49] "Neo documentation." [Online]. Available: <https://docs.neo.org/docs/en-us/basic/whitepaper.html>
- [50] F. Muratov *et al.*, "YAC: BFT consensus algorithm for blockchain," *CoRR*, vol. abs/1809.00554, 2018. [Online]. Available: <http://arxiv.org/abs/1809.00554>
- [51] "Hyperledger iroha documentation." [Online]. Available: <https://iroha.readthedocs.io/en/latest/overview.html>
- [52] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Manubot, Tech. Rep.*, 2019.
- [53] "Poet 1.0 specification." [Online]. Available: <https://sawtooth.hyperledger.org/docs/core/releases/1.0/architecture/poet.html>
- [54] V. Costan and S. Devadas, "Intel sgx explained," *IACR Cryptology ePrint Archive*, vol. 2016, p. 86, 2016.
- [55] L. Chen *et al.*, "On security analysis of proof-of-elapsed-time (poet)," in *Stabilization, Safety, and Security of Distributed Systems*, P. Spirakis and P. Tsigas, Eds., 2017, pp. 282–297.
- [56] D. Schwartz *et al.*, "The ripple protocol consensus algorithm," *Ripple Labs Inc White Paper*, vol. 5, no. 8, 2014.
- [57] H. Yi and F. Wei, "Research on a suitable blockchain for iot platform," in *Recent Developments in Intelligent Computing, Communication and Devices*, S. Patnaik and V. Jain, Eds. Springer, 2019, pp. 1063–1072.
- [58] S. Mahankali, *Blockchain: The Untold Story: From birth of Internet to future of Blockchain*. BPB Pubs, 2019.
- [59] "Exonum documentation." [Online]. Available: <https://exonum.com/developers>
- [60] J. Carlyle, "Corda performance: To infinity...and beyond!" [Online]. Available: <https://www.r3.com/wp-content/uploads/2018/04/Corda-Performance-ENG.pdf>
- [61] "Ripple documentation." [Online]. Available: <https://ripple.com/xrp>
- [62] [Online]. Available: <https://www.hyperledger.org/category/hyperledger-explorer>
- [63] "Hyperledger iroha: Working towards standard protocols for identity and digital asset management." [Online]. Available: [https://www.jpjdec.or.jp/sp/topics/event/u71kba00000an5u-att/15\\_soramitsu.pdf](https://www.jpjdec.or.jp/sp/topics/event/u71kba00000an5u-att/15_soramitsu.pdf)