

Efficient and DoS-resistant Consensus for Permissioned Blockchains

SUBMISSION ID: 216

Existing permissioned blockchain systems designate a fixed and explicit group of committee nodes to run a consensus protocol that confirms the same sequence of blocks among all nodes. Unfortunately, when such a permissioned blockchain runs on a large scale on the Internet, these explicit committee nodes can be easily turned down by denial-of-service (DoS) or network partition attacks. Although recent studies proposed scalable BFT protocols that run on a larger number of committee nodes, these protocols' efficiency drops dramatically when only a small number of nodes are attacked.

In this paper, we propose a novel protocol named EGES that leverages Intel SGX to develop a new abstraction called "stealth committee", which effectively hides a committee into a large pool of fake committee nodes. EGES selects a different stealth committee for each block and confirms the same blocks among all nodes with overwhelming probability. Evaluation shows that EGES is the first efficient permissioned blockchain's consensus protocol, which simultaneously satisfies two important metrics: (1) EGES can tolerate tough DoS and network partition attacks; and (2) EGES achieves comparable throughput and latency as existing fastest permissioned blockchains' consensus protocols. EGES's source code is available on github.com/performance21-p216/eges.

1 INTRODUCTION

A blockchain is a distributed ledger recording transactions maintained by nodes running on a peer-to-peer (P2P) network. These nodes run a consensus protocol to ensure consistency: nodes *confirm* (i.e., agree on committing [11, 16, 36]) the same sequence of blocks (i.e., no forks). Each block contains the hash of its previous block, forming an immutable hash chain. A blockchain can be permissioned or permissionless. A typical permissionless blockchain does not manage membership for nodes and is usually equipped with a cryptocurrency mechanism (e.g., Bitcoin [63]) to incentivize nodes to follow the blockchain's protocol [42, 45].

In contrast, a permissioned blockchain runs on a set of authenticated member nodes and can leverage the mature Byzantine Fault-Tolerant (BFT) protocols [39, 75, 84] to achieve better efficiency (i.e., throughput and latency). This paper focuses on permissioned blockchains because their decoupling from cryptocurrencies has facilitated the deployment of many general data-sharing applications, including a UK medical chain [5], IBM supply chains [4], and the Libra payment system [16].

For performance and regulation reasons (e.g., meeting the honesty threshold of BFT protocols [56]), a permissioned blockchain (e.g., Hyperledger Fabric [11]) typically runs its consensus protocol on a static and explicit committee. This static committee approach is already robust for a permissioned blockchain among a small scale of enterprises [4].

Unfortunately, as permissioned blockchains are deployed on large scales on the Internet, this static committee approach is vulnerable [36, 43, 64, 79] to Denial-of-Service and network partition attacks targeting committee nodes. We discuss these two types of attacks together because any single node cannot distinguish them (§2), and we call them *targeted DoS attacks* altogether. Libra [16] also identifies DoS attacks as a significant threat but provides only partial mitigation (§2.2).

Indeed, great progress has been made in designing scalable BFT protocols (e.g., SBFT [39]) running on a larger group of committee nodes and tolerating more nodes being attacked. However, these protocols designate a small number of committee nodes to finish critical tasks (e.g., combing ACKs), making these protocols' efficiency drop dramatically if these nodes are under DoS attacks (§2.2).

With recent DoS attacks lasting for days [54, 69], tolerating such attacks is crucial, yet challenging, for applications deployed on permissioned blockchains.

To address such vulnerabilities of static committees, a promising direction is to adopt the *dynamic committee* merit from permissionless blockchain systems [20, 36]. These systems select a different

50 committee for each block (mainly for fairness) and confirm a consistent sequence of blocks with a
51 tailored consensus protocol [36, 68]. As the committee member is rotated, the dynamic committee
52 merit brings the potential to achieve liveness even if a committee is under targeted DoS attacks.

53 Simply applying the dynamic committee approach, however, is not enough for a permissioned
54 blockchain to be resistant to targeted DoS attacks. To achieve DoS-resistance, it is crucial that the
55 committee selection is unpredictable: the identities of nodes in a committee must be unpredictable to
56 the attacker before the committee tries to achieve consensus on a block. Otherwise, the attacker can
57 adaptively attack the ready-to-be committee and cause the system stuck. For instance, ByzCoin [53]
58 lets the proof-of-work winners of recent blocks be the committee, but these explicit nodes are
59 easily targeted by a DoS attacker, causing ByzCoin to lose liveness permanently [47].

60 To the best of our knowledge, Algorand is the only work that can tolerate targeted DoS attacks.
61 Algorand adopts the dynamic committee approach and lets each node independently determine
62 its committee membership based on the confirmed part of the blockchain. However, as Algo-
63 rand is designed for permissionless blockchains, it confirms a block with up to 15 rounds and
64 minute-level latency (discussed in §2.2), making it unsuitable for general data-sharing applications
65 on permissioned blockchains (e.g., Libra [16]).

66 This paper aims to explore the new design point of building a permissioned blockchain’s consen-
67 sus protocol that adopts the unpredictable dynamic committee merit to defend against targeted DoS
68 or targeted partition attacks, and at the same time, achieves comparable efficiency as existing BFT
69 protocols (e.g., SBFT [39] has second-level latency).

70 To achieve this goal, a main obstacle is to ensure that any selected committee meets the honesty
71 requirement for byzantine problems: for consistency, each committee must have at most one-
72 third of nodes being malicious [56]. Permissionless blockchains meet this requirement by selecting
73 committees based on nodes’ wealth in the cryptocurrency (i.e., proof-of-stake), but cryptocurrencies
74 are usually unavailable in permissioned blockchains. Consequently, to meet such a requirement, one
75 has to unrealistically assume that almost all member nodes (>90%) are honest (see §2).

76 Fortunately, the recent pervasive usage of SGX [26] in blockchain systems (e.g., Microsoft
77 CCF [6], REM [87], Ekiden [23], Intel PoET [66]) shows that the code integrity feature of SGX
78 can surmount this obstacle. For instance, a recent implementation [2] of MinBFT leverages SGX
79 to ensure that a node cannot send different messages to different nodes and is incorporated
80 into Hyperledger [11]. However, these systems still run their consensus protocols on a group
81 of fixed and explicit committee nodes, making them susceptible to targeted DoS attacks.

82 We present EGES¹, the first efficient consensus protocol that can tackle targeted DoS or targeted
83 partition attacks for a permissioned blockchain. EGES adopts the dynamic committee merit to
84 select a different committee for confirming each block. To defend against DoS or partition attacks
85 targeting the committees, we leverage the integrity and confidentiality features of SGX to present
86 a new abstraction called *stealth committee*.

87 EGES’s stealth committee has two new features. First, EGES selects a stealth committee in SGX: the
88 selection progress has no communication among committee nodes, and the selection result cannot be
89 predicted from outside SGX. This ensures that a committee node stays stealth (cannot be targeted by
90 the attacker) before sending out its protocol messages. Second, when nodes in a committee are trying
91 to confirm a block, EGES hides them into a large pool of fake committee nodes that behave identically
92 as the real ones observed from outside SGX, so that an attacker cannot identify the real committees.

93 However, even equipped with SGX and stealth committee, it is still challenging to efficiently
94 ensure both consistency (i.e., no two member nodes confirm conflicting blocks) and reasonable
95 liveness (i.e., allow non-empty blocks to get confirmed) in the asynchronous Internet due to the FLP
96

97 ¹EGES stands for Efficient, GEneral, and Scalable consensus.
98

| Protocol | DoS and partition resistance | With SGX? | Number of nodes | Tput (txn/s) | Confirm latency (s) |
|-------------|------------------------------|-----------|-----------------|--------------|---------------------|
| EGES | high | Yes | 300 | 3226 | 0.91 |
| | | | 10K | 2654 | 1.13 |
| Algorand | high | No | 10K | ~727 | ~22 |
| PoET | medium* | Yes | 100 | 149 | 45.2 |
| Ethereum | medium* | No | 100 | 178 | 82.3 |
| SBFT | low | No | 62 | 1523 | 1.13 |
| MinBFT | low | Yes | 64 | 2478 | 0.80 |
| BFT-SMaRt | low | No | 10 | 4512 | 0.67 |
| Tendermint | low | No | 64 | 2462 | 1.31 |
| HotStuff | low | No | 64 | 2686 | 2.63 |
| HoneyBadger | low | No | 32 | 1078 | 9.39 |

Table 1. Comparison of EGES to baseline protocols. DoS resistance is analyzed in §2; evaluation setup is in §7. EGES is the only protocol that is both DoS-resistant and is among the fastest consensus protocols for permissioned blockchains. * PoET and Ethereum cannot ensure consistency on network partition attacks [64].

impossibility [34]. Specifically, suppose a committee node x for the n^{th} block fails to receive the $(n-1)^{\text{th}}$ block after a timeout, x cannot distinguish whether it is because the committee for the $(n-1)^{\text{th}}$ block failed to confirm the $(n-1)^{\text{th}}$ block, or because x itself does not receive the confirmed $(n-1)^{\text{th}}$ block due to network problems. As the committee nodes for the $(n-1)^{\text{th}}$ block may be under DoS attacks and be unreachable, x must have a mechanism to distinguish these two scenarios in order to maintain both consistency and reasonable liveness in EGES.

EGES tackles this challenge by leveraging simple probability theory. EGES’s committee for each block contains one proposer and n_A (e.g., 300) acceptors, randomly and uniformly selected from all nodes. The proposer broadcasts its block proposal to all nodes by P2P broadcasts and seeks quorum ACKs from the acceptors. EGES models the randomly selected acceptors as a sampling of the *delivery rate* of the proposal in the P2P overlay network [36]. In the previous example, EGES confirms the proposal for the $(n-1)^{\text{th}}$ block only if the proposal is delivered to a large portion of member nodes; if multiple rounds ($D = 4$ by default) of the sampling show that very few nodes have received that proposal for the $(n-1)^{\text{th}}$ block, nodes in EGES consistently confirm the $(n-1)^{\text{th}}$ block as an empty block (with an overwhelming probability).

In sum, EGES efficiently enforces consistency and can defend against targeted DoS or partition attacks. Specifically, EGES defends against such attacks by (1) letting committee nodes stay stealth *before* they start achieving consensus for a block, (2) using fake committee nodes to conceal real committee nodes *while* they are achieving consensus for a block, and (3) switching to a different committee and consistently confirming a block even if the attacker luckily guesses most real committee nodes for this block.

In essence, EGES’s stealth committee is a moving target defense approach [24, 81] that unpredictably replaces the committee to make a DoS attacker cannot launch effective targeted attacks. EGES is efficient because confirming a block in a gracious run (e.g., the proposer can reach most acceptors) only involves two P2P broadcasts and a UDP one-way delay (§4). We provide a rigorous analysis of EGES’s DoS-resistance and proof of EGES’s consistency guarantee in §5.

We implemented EGES using the codebase from Ethereum [21] and compared EGES with nine consensus protocols for blockchain systems, including five state-of-the-art efficient BFT protocols for permissioned blockchains (BFT-SMaRt [75], SBFT [39], HoneyBadger [62], and HotStuff [84]), two SGX-powered consensus protocols for permissioned blockchains (Intel-PoET [66] and MinBFT [82]), the default consensus protocol in our codebase (Ethereum-PoW [21]), and two permissionless blockchains’ protocols that run on dynamic committees (Algorand [36] and Tendermint [20]). We ran EGES on both our cluster and AWS. The extensive evaluation results (Table 1) show that:

- EGES is robust. Among all consensus protocols for permissioned blockchains, EGES is the only protocol that can defend against targeted DoS and network partition attacks, by both a theoretical analysis (§5) and evaluation (§7.2).
- EGES is efficient. EGES confirms a block with 3000 transactions in less than two seconds in typical geo-distributed settings, comparable to evaluated consensus protocols that cannot tolerate targeted DoS attacks.
- EGES’s throughput and latency are scalable to the number of nodes. When running 10k nodes, EGES showed 2.3X higher throughput and 16.8X lower latency than Algorand with 10k nodes.

Compared to existing BFT protocols [39, 62, 75, 84] and SGX-powered consensus protocols [66, 82] for permissioned blockchains, EGES is the only protocol that can tolerate targeted DoS attacks, and EGES’s efficiency is comparable to the fastest of these protocols. Compared to Algorand, the only known DoS-resistant consensus protocol for permissionless blockchains, EGES has much higher throughput and lower latency.

Our contribution is three-fold. First, EGES leverages SGX to explore the new design point of tackling DoS attacks while enforcing both consistency and reasonable liveness (including efficiency) for a permissioned blockchain in the asynchronous Internet. Second, we designed the new stealth committee abstraction and implemented EGES’s consensus protocol. Our third contribution includes an implementation of the EGES prototype and the extensive experiments of EGES and existing blockchain consensus protocols on diverse adversarial network conditions, including targeted DoS attacks, ubiquitous DoS attacks, and network partitions. Our paper reveals that, in addition to safety and performance, DoS resistance is also an essential evaluation metric for practical Internet-scale blockchain applications (e.g., e-voting [70] and payment [16]). For instance, SGX-ToR [51], a blockchain client anonymity service, is shown [73] to be susceptible to DoS attacks targeting its directory service; deploying SGX-ToR on EGES can make SGX-ToR DoS-resistant (§7.5).

In the rest of the paper, §2 introduces EGES’s background and motivation; §3 gives an overview of EGES; §4 introduces EGES’s consensus protocol; §5 analyzes the safety and liveness of EGES; §6 covers our implementation; §7 shows our evaluation, and §8 concludes.

2 BACKGROUND AND RELATED WORK

We discuss targeted DoS and network partition attacks together because these two attacks cannot be effectively distinguished in an asynchronous network. When a node cannot reach a remote node, the node cannot determine whether it is because the remote node is under DoS attacks or because the network is partitioned. Therefore, EGES maintains consistency by handling both cases together.

Although there exist many influential systems [61, 83] addressing DDoS attacks on specific nodes, EGES is complementary to them because EGES handles diverse attack scenarios (e.g., an attacker controls nodes’ P2P modules to cause network partition [64, 79], see §3.2) from the consensus layer.

We assume that the attacker has an attack budget B (e.g., $B = 300$): the attacker can mount DoS attacks targeting B nodes at a time. We will formally define the threat model in §3.2.

2.1 Intel SGX

Intel Software Guard eXtension (SGX) [26] is a hardware feature on commodity CPUs. SGX provides a secure execution environment called an enclave, where data and code execution cannot be seen or tampered with from outside. Code outside enclaves can enter an enclave by ECalls, and SGX uses remote attestations [26] to prove that a particular piece of code is running in an enclave on a genuine SGX-enabled CPU. SGX provides a trustworthy random source (`sgx_read_rand`), which calls the hardware pseudo-random generator through the RDRAND CPU instruction seeded by on-chip entropy sources [26]. Previous studies show that this random source complies with security

197 and cryptographic standards and cannot be seen or tampered with from outside enclaves [14, 41].

198 Recent work leverages SGX to improve diverse aspects of blockchain systems. Intel’s PoET [66]
199 replaces the PoW puzzles with a trusted timer in SGX; EGES is more efficient than PoET (§7.1).
200 REM [87] uses SGX to replace the useless PoW puzzles with useful computation (e.g., big data),
201 orthogonal to EGES. Microsoft CCF [6] is a permissioned blockchain platform using SGX to achieve
202 transaction privacy, but it does not include a DoS-resistance approach. Scifer [9] uses SGX to
203 maintain reliable node identities on the blockchain, which is adopted in EGES (Appendix D).

204 Ekiden [23] and ShadowEth [85] offload the execution of smart contracts to SGX-powered
205 nodes to avoid redundant execution and to preserve privacy; TEEChain [57] uses SGX to build an
206 efficient and secure off-chain payment channel; Town Crier [86] uses SGX to build a trustworthy data
207 source for smart contracts; Tesseract [17] uses SGX to build a cross-chain coin exchange framework;
208 Obscuro [80] uses SGX to improve bitcoin’s privacy; these systems do not focus on consensus
209 protocols and are orthogonal to EGES.

2.2 Consensus for Permissioned Blockchains

210
211 We briefly introduce recent notable consensus protocols for permissioned blockchains, which are
212 also EGES’s evaluation baselines. Overall, *all* these protocols run on a static committee. To ensure
213 liveness under a DoS attacker with an attack budget of B , these protocols must scale to $3 \times B + 1$ nodes
214 (for BFT protocols) or $2 \times B + 1$ nodes (for SGX-powered protocols). However, to our best knowledge,
215 no existing protocol can achieve such scalability.
216

217 BFT-SMaRt [75] is an optimized implementation of PBFT [22]. As each node broadcasts consensus
218 messages to all other nodes, BFT-SMaRt has $O(n^2)$ message complexity to the number of committee
219 nodes, resulting in poor scalability. Its paper [75] only evaluated up to 10 nodes. SBFT [39] is a
220 scalable BFT protocol that uses a new type of committee nodes called collectors. A node sends its
221 consensus messages to only c (usually $c < 8$) explicit collectors who will then broadcast a combined
222 message using the threshold signature. SBFT’s fast path can commit a block if fewer than c nodes fail;
223 however, SBFT’s performance drops dramatically if an attacker targets the c collectors (§7.4).

224 HotStuff [84] is a BFT protocol optimized for frequent leader changes, and Libra [16] leverages
225 Hotstuff to tolerate targeted DoS attacks *on leaders*. However, since Hotstuff reports a near-linear
226 increment of latency with an increasing number of nodes, it only evaluated up to 128 nodes,
227 where an attacker can DoS attack or partition one-third of all nodes rather than finding the
228 leader. HoneyBadger [62] uses randomization to remove the partial synchrony assumption of PBFT.
229 However, both its paper and our evaluation show that HoneyBadger achieves high latency due to
230 many rounds of broadcasts in its asynchronous byzantine agreements.

231 MinBFT [82] is an SGX-powered BFT protocol with the same fault model as EGES. MinBFT
232 reduces the number of rounds in PBFT and can tolerate more faulty node failures, but MinBFT
233 still has $O(n^2)$ message complexities, so its performance is not scalable to the number of nodes.

2.3 Consensus for Permissionless Blockchains

234
235 Existing permissionless blockchains can be divided into two categories based on how they confirm
236 blocks. The first category confirms block with variants of the longest-chain rule (i.e., Nakamoto con-
237 sensus [63]), including BitCoin [63], Ethereum [21], BitCoin-NG [32], Snow-White [18], Ouroboros [50],
238 Paros [27], Genesis [15], and GHOST [74]. Specifically, each node asynchronously selects the longest
239 chain it received and confirms a block when there are k blocks succeeding it. However, waiting for k
240 more blocks leads to a long confirm latency, and previous work [35] shows that this k must be
241 large enough to ensure consistency. Moreover, the longest-chain rule cannot ensure consistency
242 under partition attacks [12, 36, 43]. Intuitively, during a network partition, each partition will
243 independently grow a chain; if these chains diverge for more than k blocks, nodes in different
244
245

partitions will confirm conflicting blocks.

The second category of permissionless blockchains confirms blocks using the committee-based BFT approach, which can confirm a block as soon as the BFT consensus is achieved. This category includes Algorand [36], ByzCoin [53], Tendermint [20], and PeerCensus [28]. These systems select distinct (dynamic) committees for different blocks based on the content (e.g., nodes' wealth) on the blockchain for fairness and for handling nodes joining or leaving. Similar to EGES, these systems run a tailored consensus protocol (BA* in Algorand [36], Tendermint [20], Tenderbake [13], and Tenderand [68]) on dynamic committees to confirm blocks.

However, these protocols cannot be ported to a permissioned blockchain because of the tight coupling with cryptocurrency. For instance, although Tendermint [10] and Tenderbake [13] are described as stand-alone BFT protocols, they assume that in any committee, fewer than one-third of nodes are malicious. In a permissioned blockchain without cryptocurrency, if we want to ensure that any randomly selected committee (say 100 nodes) from a large number of (say 10k) nodes meets this requirement with overwhelming probability ($> 1 - 10^{-10}$), we need to assume over 91% of all nodes being honest (by the hypergeometric distribution), which is an overly-strong assumption for a practical large-scale blockchain system (e.g., a global payment system [16]) on the Internet.

Moreover, these systems (except Algorand) cannot ensure liveness under targeted DoS attacks because they select committees in a predictable way so that all nodes can verify the identities of committees. For instance, ByzCoin [53] lets the proof-of-work winners of recent blocks be the committee. However, these nodes with explicit identities are easily targeted by a DoS attacker, and ByzCoin may lose liveness permanently [47] if more than one-third of these nodes are attacked. Algorand defends against targeted DoS attacks by letting each node use verifiable random functions to determine its committee membership. We provide a detailed comparison showing why EGES is more efficient than Algorand in §4.3.

3 OVERVIEW

3.1 Architecture

EGES is a consensus protocol for a permissioned blockchain running on M member nodes (*nodes* for short) connected with an asynchronous network. EGES adopts the hybrid fault model used in existing SGX-powered consensus protocols [25, 82, 88], where each node has a trusted module (i.e., the SGX enclave) that will only fail by crashing, and all other components can behave arbitrarily.

Figure 1 shows the architecture of an EGES node. Each node is equipped with an attested SGX enclave running only EGES's consensus protocol. The blockchain application layer, the transaction generation and query libraries (e.g., web3.js [8]), and the blockchain storage module are outside the enclave because they are already cryptographically protected. The P2P network and operating system are outside the enclave and untrusted.

For each block index n , EGES determines one committee among all nodes (uniqueness proved in §5.1), and a node's *committee selection module* (§4.2) knows whether the node is a committee member only within its enclave. Each committee has one *proposer* (denoted as P_n), a group of *acceptors* (denoted as A_n) with the count of n_A , and a group of *arbiters* with the count of n_α . A committee node's enclave activates corresponding *committee logic modules* for this n^{th} block according to its committee

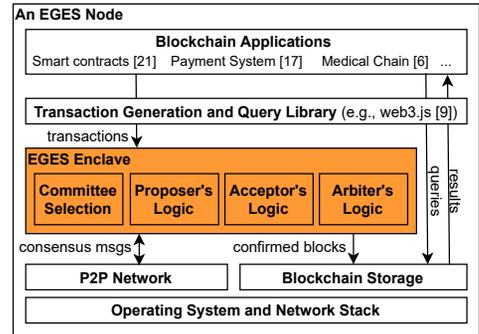


Fig. 1. An EGES node's architecture. EGES's consensus protocol has four components running in the enclave. Block proposals are included in consensus messages.

roles if the node is a committee member, and the committee logic modules will generate or respond to consensus messages following EGES’s consensus protocol (§4.3).

The proposer P_n takes a batch of transactions from outside the enclave, generates a unique block proposal (denoted as $proposal_n$) within its enclave, and tries to finalize it by collecting quorum ACKs from the acceptors (§4.3). EGES ensures that P_n ’s identity is unknown to any node’s modules outside the node’s enclave (even if the node is controlled by an attacker) before P_n broadcasts its $proposal_n$, while all acceptors A_n ’s identities are unknown to the outside throughout the whole consensus process (§4.3). To handle DoS attacks targeting P_n , EGES uses the arbiters to help P_n finalize $proposal_n$ if P_n is under attack, but the arbiters do *not* generate new block proposals (§4.4).

SGX is essential for EGES due to three main reasons. First, EGES uses SGX to regulate the behaviors of randomly selected committee nodes (§4.3); otherwise the blockchain may fork if committee nodes equivocate (i.e., sending conflicting messages to different nodes). In EGES, each node has a private key that is *only visible* within the node’s SGX enclave, and the corresponding public key works as the node account saved in all other nodes’ SGX enclaves (§4.1). A valid consensus message must carry a valid signature, proving that the message is generated in the sender node’s enclave with code integrity. By doing so, a node cannot equivocate or forge protocol messages (e.g., a proposer sending out a *finalize* message without receiving quorum ACKs).

Second, EGES leverages SGX to make its committee’s identities stealth: only a node’s enclave knows whether itself is a committee member for the current block (§4.2). This not only enables EGES to maintain practical liveness under DoS attacks, but more importantly, makes EGES’s consistency model resistant to targeted attacks. Specifically, EGES leverages probability theory to model randomly selected acceptors as a uniform sampling of the delivery rate of a block proposal in the P2P network (same as Algorand [36], see §4.3). If acceptors’ identities are public, an attacker can selectively transfer or drop packets towards them, breaking EGES’s safety.

Third, the usage of SGX enables EGES to select a stealth committee with a known count. As illustrated later in §4.3, this helps EGES to be more efficient than Algorand.

EGES has the following design goals:

- **Safety (consistency).** EGES ensures safety in an asynchronous network. Formally, if a node confirms a block b as the n^{th} block on the blockchain, the probability that another node confirms $b' \neq b$ as the n^{th} block is overwhelmingly low ($< 10^{-10}$).
- **DoS-resistance (liveness).** In addition to safety, EGES can make progress (i.e., allow non-empty blocks to be confirmed) with the assumptions about DoS attackers’ capability as described below.

Let us reconsider the subtle challenge we mentioned in §1 with a concrete example in Figure 2. This challenge is unique in EGES because EGES uses different committees for different blocks to resist DoS attacks targeting the committee. When a node cannot receive a block after a timeout, for liveness, the node cannot wait forever, but for safety, the node must figure out whether this block may have been confirmed by some nodes. Existing consensus protocols on static committees use a view change protocol that queries how many nodes have sent out ACKs and leverages quorum intersections [22, 60, 65] to address this problem. However, in EGES, this method is not viable because EGES must ensure liveness even if most nodes in A_n are DoS attacked after sending out their ACKs.

Figure 2 shows two subtle cases illustrating the challenge. Note that these two cases are deliberately simple for a clear exposition of EGES’s idea, and EGES can ensure safety in all scenarios in the asynchronous network (§5.1). In case (1), the proposer for the n^{th} block (P_n) failed before broadcasting its $proposal_n$. Therefore, $proposal_n$ is not confirmed on any node, and all nodes can safely confirm an empty block at the n^{th} position. In case (2), nodes are divided into two partitions right after the $n - 1^{th}$ block is confirmed. P_n and most nodes of A_n are in partition 1, so nodes in partition 1 can confirm the $proposal_n$ successfully. In this case, although nodes in partition 2 cannot

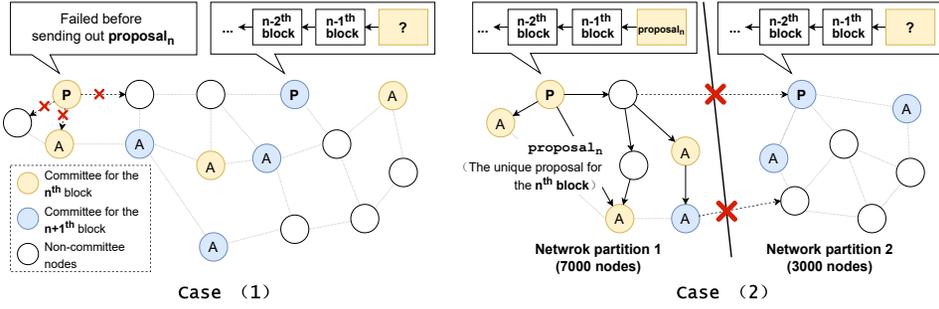


Fig. 2. A key challenge of EGES is to determine whether a block (the n^{th} block in this example) has ever been confirmed when a node cannot receive the block after a timeout. ‘P’ means the proposer; ‘A’ means an acceptor; arbiters are omitted in this figure for brevity.

receive *proposal_n*, they should *not* confirm an empty block. From any single node points of view, these two cases cannot be effectively distinguished.

To address this challenge, EGES introduces a new consensus protocol based on probability theory (§4.3). Specifically, if *proposal_n* is confirmed on some nodes, *proposal_n* should have been delivered to a large-enough portion of nodes in the P2P network because (1) confirming *proposal_n* needs quorum ACKs from nodes in \mathbb{A}_n , and (2) \mathbb{A}_n is uniformly selected from all nodes. Therefore, if we repetitively sample many nodes from all nodes, and no node has received *proposal_n*, we can predicate only a small portion of (or no) nodes have received *proposal_n*, and thus the probability of *proposal_n* having been confirmed is overwhelmingly low. To be DoS-resistant, these multiple rounds of checking must be initiated by different nodes, so EGES lets the proposers for subsequent blocks (i.e., P_{n+1} , P_{n+2} , etc.) do such samplings while seeking ACKs for their own proposals (§4.3).

3.2 Threat Model

SGX’s threat model. EGES has the same threat model for SGX as typical SGX-based systems [2, 46, 58, 67, 72]. We trust the hardware and firmware of Intel SGX, which ensures that code and data in an enclave cannot be seen or tampered with from outside. We trust that the remote attestation service can identify genuine SGX devices from fake ones (e.g., emulated with QEMU). Side-channel and access pattern attacks on SGX are out of the scope of this paper. Moreover, the adversary cannot break standard cryptographic primitives, including public-key based signatures and collision-resistant hash functions.

Transaction model. Same as most existing blockchain systems [36, 66, 78, 87], EGES assumes that (1) each transaction has a verifiable client signature, and (2) the execution and validation of any transaction are deterministic and can be performed by any node independently.

Communication model. EGES maintains safety in an asynchronous network, where network packets can be dropped, delayed, or reordered arbitrarily. Nodes may be nonresponsive, due to going offline or targeted DoS attacks (e.g., botnet DDoS attacks [54]) by a DoS attacker. When a node cannot reach a remote node, the node cannot determine whether the remote node is under DoS attack (or is offline) or the network packets are delayed. Nodes are equipped with loosely synchronous clocks (e.g., by running NTP), but EGES does not rely on the correctness of these clocks for safety.

To achieve liveness, same as existing protocols [22, 55, 75, 84]. EGES has the partial synchrony [31] assumption: there is an unknown global stabilization time (GST), after which messages between two nodes not under DoS attacks can be delivered within a known time bound.

Nodes are connected with a P2P overlay network, same as existing large-scale blockchain systems [36, 63]. Each node has a P2P module connecting to a random set of other nodes and relays messages using the gossip protocol [40, 48, 49, 71].

393 A node's P2P module is outside SGX and can be controlled by the attackers [79]: the attacker
394 can partition some nodes from other nodes [43, 64]) or selectively pass consensus messages to nodes'
395 SGX enclaves. However, such manipulations are already included in EGES's asynchronous network
396 assumption. For safety, EGES leverages the sampling merit to estimate the delivery rate of a specific
397 block proposal and derives overwhelming probability, regardless of how nodes are connected. For
398 liveness, EGES can tolerate the adversary controlling the P2P modules of a number of nodes with the
399 restriction of the adversary's attack budget described below.

400 **Assumptions on the capability of DoS attackers.** EGES has three assumptions on the capability of
401 a DoS attacker, same as Algorand [36] and existing move target defense (MTD) systems [24, 81]. First,
402 the adversary has a targeted attack budget B : the adversary cannot *constantly* cause more than
403 B targeted nodes in EGES to be nonresponsive. B can be a constant number (e.g., 300) or a fraction (e.g.,
404 10%) of the total number of nodes, but B should be bounded by the ubiquitous attack threshold in the
405 second assumption below. Note that this budget B is adaptive: the adversary can attack different
406 nodes at different times, but the number of attacked nodes at a time cannot be constantly larger than
407 B . By "constantly", we mean that the adversary can sometimes cause more than B targeted nodes to
408 be non-responsive (e.g., because most randomly selected committee nodes are luckily controlled by
409 the adversary) and cause an EGES block cannot be confirmed. However, EGES can still achieve liveness
410 by letting subsequent committee consistently confirm this block (§4.3).

411 Second, the attacker can conduct ubiquitous DoS attacks (without targeting specific nodes) or
412 partition a number of nodes from other nodes (e.g., by manipulating nodes' P2P modules [43, 64]).
413 However, the P2P overlay network should have a large enough portion (e.g., 65%) of nodes connected.
414 We provide a quantitative analysis of how EGES can preserve liveness under such attacks in §5.3.

415 Third, the adversary cannot constantly succeed in mounting an attack targeting a node within the
416 time window for the node to send out an EGES protocol message. Specifically, EGES protocol messages
417 are larger than the network maximum packet size and are fragmented into multiple packets; an
418 EGES committee node's identity is unknown to the adversary before sending out the first packet,
419 and we assume that the adversary cannot mount targeted DoS attacks until the node sends out
420 all packets belonging to this message (at most hundreds of kB and can be sent within one second).

421 EGES already assumes a strong enough attacker for practical distributed systems on the Internet.
422 As pointed out by Algorand [36], a more powerful adversary than our model usually controls
423 the internet service provider and can prevent all EGES nodes from communicating at all: no
424 practical system can ensure liveness under such a strong adversary, and such attacks can be
425 easily detected. We will provide a rigorous analysis of EGES's DoS resistance in §5.2.

426 **DoS resistance of EGES.** EGES has three important features to achieve DoS resistance:

- 427 • EGES randomly selects a distinct committee for each block. The selection is done inside the SGX
428 enclaves of a previous committee, and the selection result is encrypted on the confirmed common
429 prefix of the blockchain. By doing so, a committee node can determine its committee membership
430 without interactions with other nodes, making it *stay stealth* before trying to achieve consensus on
431 its block.
- 432 • When a committee is achieving consensus for a given block, EGES uses fake committee nodes
433 to conceal the real ones by sending encrypted dummy messages. Since whether a node is a
434 real committee node is only known within the node's SGX enclave, and the encrypted dummy
435 messages are of the same format as real ones, a DoS attacker cannot distinguish the real committee
436 nodes from the fake ones. Therefore, the attacker must have an unrealistic large attack budget to
437 attack all the real and fake committees; otherwise, he has to randomly guess who are the real ones.
- 438 • Even if the attacker luckily guesses the real ones (he may eventually succeed if trying persistently),
439 EGES can ensure safety with overwhelming probability. Specifically, even if a committee cannot
440

confirm its own block, committees for subsequent blocks can help to consistently confirm this block (§4.3). This feature is in contrast to most existing consensus protocols (i.e., all except Algorand [36]), where the system must wait statically until a quorum of nodes become reachable.

4 EGES CONSENSUS PROTOCOL

4.1 Protocol Preliminaries

Protocol parameters. EGES’s consensus protocol has three parameters, n_A (default 300), τ (default 59%), and D (default 4), where n_A is the number of acceptors, τ is the quorum ratio, and D is the finalization depth for an empty block. We will show how to select these parameters in §5.3 and how these parameters affect EGES’s performance in §7.3.

Block structure. EGES adds one data field to the block structure of common blockchain systems [21, 63]: the encrypted committee identities for a future block (§4.2). EGES is oblivious to how transactions are stored or executed.

Invariant 1 (see §5.1 for proof). *For any block index n , at most one unique block proposal ($proposal_n$) is generated; a node can only confirm $proposal_n$ or a default empty block ($empty_n$) as its n^{th} block.*

Block status. Each block in a node’s chain has three states: *undecided*, *finalized*, and *confirmed*. An undecided n^{th} block can only be $empty_n$. A node appends $empty_n$ to its chain when the node triggers a timeout waiting for the finalize message for the n^{th} block; the block is in the undecided state because the node cannot determine (for now) whether it should confirm $proposal_n$ or $empty_n$.

A finalized n^{th} block can be either $proposal_n$ or $empty_n$. EGES ensures the following invariant:

Invariant 2. *If a node’s n^{th} block is finalized as $proposal_n$, no other nodes will finalize the n^{th} block as $empty_n$, and vice versa.*

There are two rules for appending a finalized n^{th} block: (1) a node appends finalized $proposal_n$ if it receives the finalize message for $proposal_n$, and (2) a node changes the $empty_n$ from the undecided state to finalized if the node can predicate that no node has finalized $proposal_n$ (§4.3).

A node confirms its finalized n^{th} block if all blocks with indices smaller than n in its chain are finalized. Note that although EGES may finalize blocks out of order, EGES confirms blocks sequentially, same as typical blockchains [11, 21].

Each node’s local states. Each EGES node maintains three major local states: a local blockchain (the chain), a proposal cache (the cache), and a set of `learnedProposals`. The cache is maintained in the node’s EGES enclave. When the node receives $proposal_n$, it puts the proposal into the cache, in case the committees of future blocks query the delivery rate of $proposal_n$.

The chain on each node is divided into two parts: the confirmed part and the unconfirmed part. We use `MC` to represent the maximum confirmed index and `U` to represent the indices of undecided blocks in chain. The confirmed part of chain (i.e., indices $\leq MC$) are cryptographically-chained by hash values and can be saved out of the enclave and get executed, while unconfirmed parts are saved in the EGES enclave. The `learnedProposals` is the set of known proposals for undecided blocks on this node and is saved in EGES enclave.

Membership and key management. Each node i has a key pair $\langle pk_i, sk_i \rangle$, with the public key pk_i as its account, and its secret key sk_i is *only visible within its enclave*: even this node’s administrator cannot see the plain-text of its secret key. We use the notations from PBFT [22]: we denote a message m_1 sent to node i encrypted by i ’s public key pk_i as $\{m_1\}_{pk_i}$; we denote a message m_2 generated by node i ’s enclave and signed by sk_i as $\langle m_2 \rangle_{\sigma_i}$. For efficiency, EGES signs on message digests.

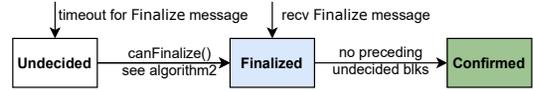


Fig. 3. EGES’s block status diagram.

491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

| Propose | |
|-----------------------------------|--|
| n | The index of the proposal |
| blk | The content of the block to be proposed |
| MC _p | The MC value of the proposer |
| U _p | The U list of the proposer |
| proposal _{u_m} | $u_m = \max(U_p)$. The proposer tries to finalize this proposal together with its proposed n^{th} block |
| ACK | |
| n | The index of corresponding proposal |
| sender | The sender's public key (account) |
| notifications | Proposals to notify the proposer |
| isReal | For identifying real ACK from cover messages |
| nonce | Random padding to make cipher text unpredictable |
| Finalize | |
| n | The index of the block finalized |
| u _m | The index of the block that is finalized together |
| learnt | Proposals learnt from acceptors notifications |

Table 2. EGES's messages' fields. Blue fields are used only in the checking mode (nil in the normal mode).

Algorithm 1: the proposer for the n^{th} block

```

1  $n_A \leftarrow$  the number of acceptors
2 blk  $\leftarrow$  the content of the  $n^{th}$  block to propose
3 function normalPropose():
4   bcast (Propose, blk, MC, nil, nil) $\sigma_{me}$ 
5   upon receiving (ACK, n, pki, nil) $\sigma_i$ 
6     ACKs.insert(pki)
7     if ACKs.count  $\geq$   $\tau \times n_A$ : bcast (Finalize, n, nil, nil) $\sigma_{me}$ 
8 function checkPropose():
9    $u_m \leftarrow \max(U)$ 
10  if cache[um] != nil || learntProposals[um] != nil:
11    proposalum  $\leftarrow$  the proposal for index um
12  else: proposalum  $\leftarrow$  nil
13  bcast (Propose, blk, MC, U, proposalum) $\sigma_{me}$ 
14  learnt = []
15  upon receiving (ACK, n, pki, notifications) $\sigma_i$ 
16    ACKs.insert(pki)
17    learnt.insert(notifications)
18    if ACKs.count  $\geq$   $\tau \times n_A$ :
19      bcast (Finalize, n, proposalum, learnt)
```

Algorithm 2: all nodes' action

```

/* All message senders' memberships and signatures are verified,
omitted in all algorithms for brevity */
1 upon receiving msg = (Propose, n, blk, MCp ...) $\sigma_i$ 
2 cache[n] = msg
3 if MC < MCp: ask peers for missing blocks
4 if hash(ski, n) > threshold: Reply fake (cover) ACKs message
/* same format as true ACKs, with isReal = false. */
5 upon receiving (Finalize, n, um, learnt) $\sigma_i$ 
6 if um != nil:
7   chain[um].status  $\leftarrow$  finalized
8   U = U  $\cup$  um
9 if |U| == 0: Confirm chain up to index n (MC  $\leftarrow$  n)
10 else:
11   learntProposals.insert(learnt)
12   for u in U in descending order:
13     if canFinalize(u):
14       chain[u].status = finalized
15     else: break // Empty blocks are finalized in descending order.
16 timeout waiting for next block
17 chain.append(empty, status = undecided)
18 function canFinalize(u):
19   count = 0
20   for i = u + 1; i  $\leq$  chain.len; ++i:
21     if chain[i] != empty:
22       count++
23     if count  $\geq$  D: return true
24   else: return false
```

Algorithm 3: an acceptor for the n^{th} block

```

1 upon receiving (Propose, blk, MCp, Up, proposalum) $\sigma_i$ 
2 r = rand()
3 if MCp < MC: notify proposer to catch up (omitted)
4 if |Up| == 0:
5   reply ({ACK, n, pkme, nil, true, r}) $\sigma_{me}$ 
// Only the leader(i) can decrypt this message
6 else:
7   notifications = []
8   for u in Up:
9     if cache[u] != nil: notifications.append(cache[u])
10  reply ({ACK, n, notifications, true, r}) $\sigma_{me}$ 
```

Fig. 4. EGES's consensus protocol.

To ease understanding, we describe EGES protocol on a fixed membership, where all nodes' accounts (public keys) are loaded to nodes' EGES enclaves a priori, and all nodes' EGES enclaves are attested. We show how EGES supports dynamic membership and attestations in [Appendix D](#).

4.2 Selecting a Stealth Committee

For each block, EGES selects a committee, including **one proposer**, n_A **acceptors**, and n_α **arbiters**, in an unpredictable way without communication among nodes.

The committee members for the n^{th} block is selected in the EGES enclave of P_{n-lb} , and these committee nodes' identities are encrypted in the $(n-lb)^{th}$ block. lb (look-back) is a system parameter and needs to be large enough (e.g., the number of blocks confirmed in days) to ensure that even when the network condition is poor, and new blocks cannot be confirmed in time, EGES can still derive committees for future blocks. We assume that the first lb committees' identities are encrypted in the genesis (0^{th}) block by a trusted party, or that the blockchain is bootstrapped in a controlled domain for at least lb blocks. Note that the value of lb does not affect EGES's safety.

Occasionally, a node may be selected as the committee for a future block and then leave the system, which EGES already tolerates as a failed node. If the $(n - lb)^{th}$ block happens to be an empty block, EGES uses the committee identities encrypted in the $(n - 2lb)^{th}$ block (and identities in the $(n - 3lb)^{th}$ block if the $(n - 2lb)^{th}$ block is also empty, recursively). Although this proposer's identity is already explicit when confirming the $(n - lb)^{th}$ block and may be targeted, EGES can tolerate it as a failed proposer and uses subsequent committees to confirm $empty_n$.

$P_{(n-lb)}$ selects the committee for the n^{th} block with two steps, which are done in $P_{(n-lb)}$'s EGES enclave to ensure both integrity (i.e., an attacker cannot control the selection) and confidentiality (i.e., an attacker cannot know the selection result). In the first step, $P_{(n-lb)}$ randomly selects the committee members from all member nodes following the uniform distribution. Recall that the member list is loaded in the EGES enclave on each node (§4.1), so $P_{(n-lb)}$ simply selects $n_A + 1$ nodes from the list using the SGX's trustworthy pseudo-random number generator as the random source, which has been shown to be cryptographically-secure and cannot be seen or tampered with from outside enclave (§2.1).

In the second step, for each selected committee node, $P_{(n-lb)}$ generates one certificate, which is the cipher-text of the concatenation of a predefined byte string and a random nonce (for making the cipher-text unpredictable), encrypted with that committee node's public key. Then, $P_{(n-lb)}$ includes these $(n_A + n_\alpha + 1)$ certificates in the $(n - lb)^{th}$ block's proposal. The first certificate is for the proposer, and the other n_A certificates are for acceptors. When a node confirms this block, it tries to decrypt one of these certificates using its own secret key in its enclave; if the node can get the predefined string, it predicates that it is a committee node for the n^{th} block.

Despite using asymmetric cryptography, this mechanism is efficient in EGES because both encryption and decryption are done asynchronously off the consensus's critical path. For encryption, since $P_{(n-lb)}$'s enclave knows it is the proposer for the $(n - lb)^{th}$ block after confirming the $(n - 2lb)^{th}$ block, $P_{(n-lb)}$ starts selecting the committees and encrypting the certificates as soon as confirming the $(n - 2lb)^{th}$ block. Similarly, the decryption is also off the critical path as the decryption result is used lb blocks later.

EGES's committee selection mechanism is unpredictable and non-interactive because: (1) the random source cannot be seen or tampered with from outside the enclave of P_{n-lb} , and the certificates can only be verified within a selected committee node's enclave; and (2) the selection process is solely done within the EGES enclave of P_{n-lb} . These two features ensure the committee nodes' identities are not exposed during the selection, so the committee nodes cannot be targeted before sending out protocol messages for the n^{th} block.

Discussions. EGES selects only one proposer for each block to achieve good efficiency: EGES only needs to achieve a binary consensus on whether to confirm the unique proposal by this proposer or a default empty block. For acceptors, an alternative design is to let each node *independently* determine whether it is an acceptor for the current block with a probability, and EGES only controls the *expected* total count. However, this alternative design will lead to a much larger quorum ratio (i.e., τ) to ensure safety and thus worse liveness (quantitative analysis in §7.3).

4.3 Confirming a Block

A proposer P_n has two operation modes: *normal* mode and *checking* mode. P_n is in the normal mode if all blocks in its chain before n are confirmed (i.e., $\mathbb{U} = \emptyset$), and P_n tries to confirm $proposal_n$ quickly. Otherwise, P_n is in the checking mode: while proposing $proposal_n$, it also checks the status of the undecided blocks in its chain.

Normal mode. Algorithm 1 Line 3~7 shows how a normal mode P_n tries to confirm $proposal_n$ in a gracious run. First, P_n broadcasts a propose request through the P2P network carrying $proposal_n$

589 and its MC (§4.1). The MC value helps nodes align confirmed parts of their chain: if a node’s MC
 590 is smaller than the proposer’s, the node asks for the missing confirmed blocks from its peers (Algo-
 591 rithm 2 Line 3). Upon receiving this propose request, an acceptor replies an ACK using UDP directly
 592 to P_n (Algorithm 3 Line 5). Second, P_n waits for quorum ($\tau \times n_A$) ACKs from \mathbb{A}_n . P_n does not know
 593 which nodes are acceptors, but EGES’s ensures that a non-acceptor cannot send valid ACKs (§4.1).
 594 Third, P_n broadcasts a `finalize` message; on receiving this message, a node finalizes `proposal_n`.
 595 **Checking mode.** P_n is in the checking mode if it has undecided blocks (i.e., \mathbb{U} is non-empty), and its
 596 workflow is shown in Algorithm 1 Line 8~19. P_n checks the status of its undecided blocks and tries to
 597 finalize them (if possible) by adding additional fields to the propose message.

598 Each $u \in \mathbb{U}$ in P_n ’s chain is categorized into one of the two types: (1) if P_n has learnt the unique
 599 `proposal_u`, either from the propose messages or from P_u from the notifications of other nodes, we
 600 call u a “known undecided” block; (2) otherwise we call u “unknown undecided”. For each unknown
 601 undecided block u , P_n tries to learn `proposal_u` from \mathbb{A}_n . If P_n learns the proposal, P_n carries it in the
 602 `finalize` message in order to let subsequent proposers finalize it. Otherwise, P_n carries a message
 603 stating that most acceptors in \mathbb{A}_n never received `proposal_n`, and a node receiving this message
 604 finalizes `empty_n` if the node received such messages from more than D consecutive proposers
 605 (Algorithm 2 Line 18).

606 For known undecided blocks, P_n helps to finalize *only* `proposal_{u_m}` where $u_m = \max(\mathbb{U})$ and
 607 leaves other blocks for subsequent proposers. In other words, undecided blocks must be finalized in
 608 descending order. This is because, without this restriction, when a node finalizes `empty_n`, it only
 609 ensures proposers for blocks with index $\leq n+D$ has not finalized `proposal_n`; adding this restriction
 610 helps to ensure that proposers for blocks with index $> n+D$ cannot finalize `proposal_n`.

611 To help understanding, we give three concrete examples in Appendix C showing (1) how a
 612 proposer helps to finalize an undecided proposal, (2) how a proposer finalizes an undecided block as
 613 empty, and (3) why a checking mode proposer can only finalize `proposal_{u_m}` where $u_m = \max(\mathbb{U})$.

614 **Discussions.** EGES does not let the acceptors validate transactions enclosed in `proposal_n` because
 615 EGES may finalize blocks out of order. Specifically, if `proposal_n` contains a money transfer trans-
 616 action, and an acceptor has undecided blocks preceding n , the acceptor cannot predicate whether
 617 the source account has enough balance without a complete history. Therefore, EGES validates transac-
 618 tions when they are confirmed at each node because EGES always confirms blocks in order. An invalid
 619 transaction is deterministically (§3.2) discarded by all nodes, without impairing EGES’s safety.

620 For the same reason, EGES does not force P_n to generate its `proposal_n` within its SGX enclave
 621 because P_n cannot forge a valid signed transaction to affect EGES’s safety. Overall, EGES achieve
 622 consensus on `proposal_n` opaquely.

623 The drawback is that EGES may waste resources on achieving consensus on invalid transactions.
 624 However, as all invalid transactions are recorded on the blockchain, if an entity keeps injecting
 625 invalid transactions, it can be easily detected and penalized in a permissioned deployment.

626 **Comparison with Algorand.** EGES and Algorand both select a distinct committee for each block in
 627 an unpredictable way, and both use the delivery rate of a block proposal to confirm a block. Algorand
 628 leverages its built-in cryptocurrency to incentivize committee nodes to follow its protocol (i.e., proof
 629 of stake). However, even if one runs Algorand within SGX in a permissioned blockchain, there are
 630 still two major design differences making EGES more efficient than Algorand.

631 First, Algorand uses verifiable random functions (VRF) to determine committees, so it can only
 632 control the expected count of proposers for each block without an exact number (1 ~ 70 in their ex-
 633 periment [36]). This design makes Algorand’s consensus protocol not responsive [16, 84]: informally,
 634 a responsive protocol lets nodes wait for a number of messages rather than a large amount of time
 635 in each protocol step, which ensures a good performance when the network is in good condition. For
 636 each block, Algorand selects 1 ~ 70 proposers, and each proposer broadcasts a block proposal with
 637

638 a distinct priority level. Then, Algorand selects one of these proposals by letting nodes vote for the re-
 639 ceived proposal with the highest priority. Since the total number of proposers is unknown, each node
 640 must wait for a conservatively long time (e.g., 10s) before voting to ensure it received most proposals.
 641 In contrast, EGES selects one proposer for each block, without the necessity for the selection progress,
 642 and EGES’s protocol is responsive (§4).

643 Second, EGES adopts an optimistic design while Algorand adopts a pessimistic design. Specifically,
 644 Algorand uses a heavy step for both confirming a non-empty block and confirming an empty block.
 645 In contrast, EGES optimistically makes its gracious runs (i.e., confirming proposal_n) fast and shifts
 646 the burden of maintaining consistency to the rare failure cases (i.e., confirming empty_n).

648 4.4 Handling DoS Attacks Targeting Proposers

649 In EGES, a proposer stays stealth before proposing its block, but its identity becomes explicit after
 650 broadcasting its proposal. If the proposer is DoS attacked at this time, this block cannot be finalized in
 651 time, impairing EGES’s liveness.

652 To address this problem, we propose a new role of nodes called *arbiter*. An arbiter for a
 653 block index n does *not* generate new block proposals but only helps the proposer to finalize
 654 its proposal. For each block index n , the count of arbiters n_α is large than the attack budget
 655 B , and these arbiters do the same tasks to tolerate DoS attacks targeting them.

656 On receiving proposal_n , an arbiter for the n^{th} block broadcasts an *arbit* request following
 657 the same protocol as the proposer (Algorithm 1), and an acceptor responds to an *arbit* message
 658 with the same logic responding to a *propose* message.

659 **Discussions.** With the arbiters’ help, a proposer’s critical task is only to *send out* its *propose*
 660 request, and the arbiters help to finalize it. However, responding to both proposer and multiple
 661 arbiters makes acceptors targets of DoS attacks. Therefore, EGES lets normal nodes also randomly
 662 send fake (dummy) ACKs to cover the real acceptors (Algorithm 2 Line 4). Since real or fake ACKs are
 663 all encrypted with the receiver’s public key, only the receiver’s EGES enclave can decrypt them in the
 664 enclave and distinguish the real ones, so an attacker cannot know who are the real acceptors.

666 5 SECURITY ANALYSIS

667 5.1 Safety with Overwhelming Probability

668 EGES ensures safety with overwhelming probability (i.e., $> 1 - 10^{-10}$). Formally, if a node confirms a
 669 block b as the i^{th} block on the blockchain, the probability that another member node confirms $b' \neq b$
 670 as the i^{th} block is $< 10^{-10}$.

671 We prove the safety guarantee of EGES by induction, which is shown in Appendix A.

674 5.2 Liveness under Targeted DoS Attacks

675 EGES can defend against DoS attacks and partition attacks targeting committee nodes under the
 676 threat model in §3.2. Since from a single node’s point of view, it cannot distinguish whether a remote
 677 node is under DoS attack or is partitioned, EGES handles these two attacks altogether.

678 EGES has three types of committee nodes for each block, a proposer, n_A acceptors, and n_α arbiters,
 679 randomly selected from all nodes in the system. Because of EGES’s stealth committee abstraction
 680 (§4.2), the identity of each committee node is unknown to attackers outside SGX enclaves before the
 681 node sending out its first protocol message.

682 We first discuss acceptors. An acceptor sends ACK messages to both the proposer and arbiters,
 683 and its identity becomes explicit after sending out its first ACK message, so EGES uses fake acceptors
 684 to conceal the real ones. If observed from outside enclaves, the fake acceptors behave identically
 685 as real ones so an attacker cannot differentiate the real acceptors and attack them. EGES achieves
 686

687 this with three design points. First, fake acceptors are randomly selected from all nodes for each
 688 block so that an attacker cannot determine the fake acceptors by monitoring network packets (§4.2).
 689 Second, real and fake acceptors' EGES enclaves respond to protocol messages (propose and arbit)
 690 in the same way if observed outside the enclaves (§4.3), so an attacker cannot distinguish real or fake
 691 acceptors by watching their behaviors. Third, all messages from real or fake acceptors have the same
 692 format, encrypted with the receiver's public key (§4.3) and can only be decrypted in the receiver's
 693 enclave, so an attacker cannot differentiate real acceptors from fake ones by watching the packet content.

694 Therefore, if an attacker targeting acceptors in EGES, it can only randomly select B nodes
 695 from both real acceptors and fake acceptors. For instance, if EGES has $n_A = 300$, $\tau = 59\%$ and
 696 has (expected) 600 fake acceptors; if the attacker's attack budget is 300, the probability that
 697 the attacker can luckily attack more than $(1 - \tau) \times n_A$ real acceptors for one block is 0.3%.
 698 Moreover, even if the attacker is so lucky that it successfully guessed more than $(1 - \tau) \times n_A$
 699 acceptors for some block (say n^{th}), EGES can still consistently determine whether to confirm
 700 `proposaln` or `emptyn` using the committees for subsequent blocks (§4.3); in other words, the
 701 attacker must constantly be so lucky to make EGES lose liveness.

702 Then we discuss proposers and arbiters. The identity of a proposer P_n becomes explicit after
 703 broadcasting its `proposaln` and may be targeted attacked when it is waiting for quorum ACKs.
 704 However, since EGES has many (i.e., $n_\alpha > B$) arbiters that can help to finalize the `proposaln`,
 705 attacking P_n will not affect EGES's liveness. Note that as long as one arbiter is not under targeted
 706 DoS attack, it can finalize the current n^{th} block.

707 Note that EGES's targeted attack model (§3.2) handles only DoS or partition attacks targeting
 708 specific EGES nodes. A more powerful attacker may also target EGES's major communication links.
 709 From the protocol aspect, EGES avoids such vulnerabilities in the Network layer (in the OSI model [3])
 710 by using distinct committees for different blocks: EGES's protocol traffic is spread among the whole
 711 P2P network rather than centralized among a few dedicated nodes. However, when EGES is deployed,
 712 EGES's communication messages may be aggregated in the Link or Physical layer. For instance, if a
 713 large number of EGES nodes are hosted in the same data center (DC), the links connecting this DC and
 714 the Internet may be susceptible to attacks. Fortunately, such attacks are not adaptive, and as long as
 715 a great majority of nodes are connected, EGES can achieve practical liveness. §7.3 shows the relation
 716 between EGES's liveness and the maximum connected component size in the P2P network. Never-
 717 theless, EGES cannot ensure liveness under arbitrary partitions, and previous work shows that it is impos-
 718 sible to ensure both consistency and liveness under partitions [34, 37].

720 5.3 Parameter Selection

721 EGES has two correlated parameters τ and D . We show in Appendix B the relation between τ and D
 722 to ensure EGES's safety under any network condition and to achieve reasonable liveness (confirm
 723 non-empty blocks) on network partitions (or ubiquitous DoS attacks).

725 6 IMPLEMENTATION

726 We selected the Golang implementation of Ethereum (i.e., geth) as our codebase because geth is
 727 heavily tested on the Internet. We leveraged the P2P libraries from geth and rewrote the functions
 728 for generating, verifying, and handling new blocks. Since SGX only provides SDKs in C/C++, we used
 729 CGo to invoke ECalls. We modified 2073 lines of Golang code and implemented the consensus
 730 protocol for 1943 lines of C code. For asymmetric key based encryption, we used ECC-256 from
 731 the API provided by the SGX SDK. For timeouts, we used the trusted timer API `sgx_get_trusted_time`
 732 provided by the SGX platform

733 Each EGES node has three modules: a consensus module running the EGES consensus protocol and
 734 storing nodes' member list (§4), a P2P module connecting to a random set of peers and relaying

735

| Config | # Nodes | Acceptor group size (n_A) | D | τ | LB | timeout | SGX mode |
|-----------|-----------|-------------------------------|---|--------|-------|---------|-----------------|
| Cluster | 300 | 100 | 4 | 65% | 5000 | 2s | hardware mode |
| AWS Cloud | up to 10K | 300 | 4 | 59% | 10000 | 3s | simulation mode |

Table 3. EGES’s evaluation parameters.

messages using the Gossip [49] protocol, and a blockchain core module storing confirmed parts of chain and client transactions. Only the consensus module runs in the node’s SGX enclave.

In EGES, a node may finalize a block before knowing its preceding blocks. Therefore, when an EGES proposer proposes a block or a node finalizes a block, it leaves the block’s field of “hash of the previous block” empty, and EGES’s enclave computes this field when confirming the block. In essence, EGES achieves consensus on a totally ordered sequence of transactions, same as Hyperledger Fabric [11], and encapsulates these batches into a hash-chain of blocks while confirming them.

6.1 Membership, Attestation, and Enclave Interactions

Appendix D covers how EGES supports dynamic membership, and Appendix E shows EGES’s enclave interactions and how EGES handles enclave forking attacks.

7 EVALUATION

Evaluation setup. Our evaluation was done on both our own cluster with 30 machines and the AWS cloud, with parameters shown in Table 3. In our cluster, each machine has 40Gbps NIC, 2.60GHz Intel E3-1280 V6 CPU with SGX, 64GB memory, and 1TB SSD. On AWS, we started up to 100 c5.18xlarge instances (VMs) running in the same city, each of which has 72 cores, 128GB memory, and up to 25 Gbps NIC. We ran up to 100 EGES nodes on each VM (10k nodes in total), with each EGES node running in a docker container.

To evaluate EGES and baseline protocols in a geo-replicated setting, while running EGES on both our cluster and AWS, we emulated the world scale Internet by using the Linux traffic control (TC) to limit the RTT between every two nodes to a random value between 150ms and 300ms. These settings are comparable to Algorand’s setting on AWS. As AWS does not provide SGX hardware, we ran EGES in the SGX simulation mode on AWS and in the SGX hardware mode on our cluster; we show that EGES’s performance in simulation mode is roughly the same as hardware mode because EGES’s performance is bound to network latency in WAN (§7.1). The scalability (Figure 5) and robustness (Figure 6) experiments were done on AWS, and the rest were in our cluster.

We evaluated EGES with nine consensus protocols for blockchain systems, including five state-of-the-art efficient BFT protocols for permissioned blockchains (BFT-SMaRt [75], SBFT [39], HoneyBadger [62], and HotStuff [84]), two SGX-powered consensus protocols for permissioned blockchains (Intel-PoET [66] and MinBFT [82]), the default consensus protocol in our codebase (Ethereum-PoW [21]), and two permissionless blockchains’ protocol that runs on dynamic committees (Algorand [36] and Tendermint [20]). A detailed description of these protocols is in §2.2.

Since Algorand’s open-source code is under development, and we were unable to deploy its latest release [1] to the same scale as EGES and Algorand’s paper (i.e., 10k nodes), we took Algorand’s performance from Figure 5 in its paper [36]. To make the comparison fair, we make EGES’s network setting more rigorous than Algorand’s: Algorand divided nodes into multiple cities where intra-city packets have negligible latency, while EGES lets the RTT among any two nodes be at least 150ms.

For all evaluated protocols, we measured their performance when each of them reached peak throughput. For an apple-to-apple comparison of latency, we adopted Algorand’s method to measure transactions’ server-side confirmation time: from the time a transaction is first proposed by a committee node to the time the transaction is confirmed at this node, excluding

the time for clients’ transaction submissions. We measured the server-side instead of the client-side latency because this method precludes the disturbance of client behaviors as these protocols run on different blockchain frameworks. For instance, in Ethereum, PoET (running on Hyperledger Sawtooth [66]), and EGES, a client submits a transaction to a random node, and the transaction is disseminated via P2P networks; in BFT-SMaRt, a client submits transactions to a fixed node (i.e., the leader); in Algorand, a consensus node directly packs a block with a fixed amount of data (e.g., 1MB) instead of using separate transactions.

We set EGES’s transaction size to 250 bytes, a typical transaction size for general data-sharing applications [7, 10]. Since Algorand reported throughput on block size, we convert it to txn/s by assuming the same size of transactions as EGES’s. The transaction sizes for the other eight baseline protocols are either equal to or smaller than that of EGES. Our evaluation focuses on these questions:

§7.1 : Is EGES efficient and scalable?

§7.2 : What is EGES’s performance under DoS attacks?

§7.3 : How sensitive is EGES to its parameters?

§7.4 : How do EGES performance and fault tolerance compare with notable BFT protocols?

§7.5 : What are the limitations and future work of EGES?

7.1 Efficiency and Scalability

Table 1 shows the performance comparison of EGES and eight baseline protocols. As Algorand’s paper [36] evaluated at least 2K nodes, we postpone the comparison between EGES and Algorand to when we evaluated EGES’s scalability.

Overall, in the geo-replicated setting, EGES achieved comparable performance to MinBFT, Tendermint, HotStuff, and SBFT. We ran BFT-SMaRt in its default setting (ten nodes), and it showed higher throughput and lower latency than EGES. BFT-SMaRt is more suitable for small scale permissioned blockchains where a few companies run nodes in a controlled environment, so it lets nodes send messages to each other directly. In contrast, EGES is designed for tolerating targeted DoS attacks on committee nodes, so it has two P2P broadcasts to confirm a block. §7.4 shows that EGES’s scalability and fault tolerance are better than BFT-SMaRt.

SBFT and HotStuff had a lower throughput and a higher latency than EGES. They rely on designated nodes to collect the consensus messages that were originally all-to-all broadcasted and to distribute a combined message to all nodes. Although this approach improves scalability, it also incurs two more RTTs, limiting their performance in a geo-distributed deployment. Moreover, an attacker targeting these designated nodes will cause a dramatic performance drop to the system, which is evaluated in §7.4.

HoneyBadger showed a lower throughput and a higher latency than EGES because HoneyBadger uses multiple rounds of broadcasts for a single block, which incurred a long latency in a geo-distributed setting. EGES showed orders of magnitude better performance than PoET and Ethereum, two PoW protocols. Their performance is limited by the time for solving PoW puzzles (or sleep time) and the number of blocks to wait for before confirming a block (§2.3). Our evaluation result for PoET is similar to a recent study [29].

Breakdown and micro-events. To understand EGES’s latency, we recorded the time taken for the two steps of EGES’s protocol (§4.3): seeking for quorum ACKs took 576ms; broadcasting finalize messages took 329ms. The first step took a longer time because EGES broadcasts the proposed block in its P2P network in this step. This P2P broadcast time is essential in any blockchain system because new blocks need to be broadcasted to all nodes.

SGX’s overhead. Table 4 shows the micro-events of EGES. The ECall column shows the number of times that EGES’s proposer node entered its SGX enclaves on finalizing a block. Since each ECall only

| blk size | txns/blk | # ECalls | CPU usage | network usage |
|----------|----------|----------|-----------|---------------|
| 750 KB | 3000 | 97 | 12.4% | 15.53 Mbps |

Table 4. Proposer’s micro-events for finalizing a block.

takes around 3us [26], and EGES’s proposer only did 97 ECalls on average for each block, running in SGX hardware mode and simulation mode makes little difference for EGES’s performance.

Scalability. To evaluate EGES’s scalability, we ran 100-10000 nodes on AWS and evaluated its confirm latency with the same block size as in the cluster evaluation. Figure 5 shows the result. The latency is divided into two parts. The figure shows that the seeking for quorum ACKs phase of EGES (§4.3) is the dominant factor because it broadcasts the proposed 750KB block on the P2P network. Fortunately, a P2P broadcast latency is proportional to approximately the \log of the number of nodes [76], indicating EGES’s reasonable scalability. The increase rate was slightly greater than the log scale because 100 nodes were run in one VM with CPU and NIC contentions. EGES’s latency on AWS was slightly faster than on our cluster, because AWS CPUs are faster.

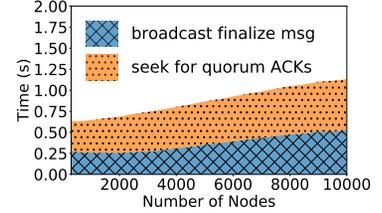


Fig. 5. Scalability to the number of nodes on the Internet.

Compared to Algorand’s performance in Table 1, EGES showed 2.3X higher throughput and 16.8X faster latency than Algorand. This is due to two reasons. First, Algorand’s VRF-based method selects multiple proposers for each block, and Algorand uses a reduction step to select one proposal by these proposers. Moreover, as the VRF-based approach cannot control the exact number of proposers, nodes must wait for a conservatively long time in the reduction step (§4.3). In contrast, EGES’s stealth committee abstraction selects one proposer for each block, without the need for such a reduction step. Second, EGES’s consensus protocol has only two rounds in gracious runs to confirm a block (§4.3).

7.2 Performance on DoS Attacks

To evaluate EGES’s robustness under DoS attacks, we ran EGES with 1000 nodes on AWS with $n_A = n_\alpha = 100$, and conducted targeted DoS attacks that are compliant with our attack model (§3.2): we assumed the attacker’s budget $B = 10\%$ of total nodes, and we set the expected count for fake acceptors and arbiters to be 200. Each time we targeted the current proposer and 99 arbiters or real/fake acceptors (because we cannot distinguish real acceptors). For each attack, we blocked all communication from the attacked nodes for 20 seconds.

We deem such attacks to be powerful enough, as no existing protocols for permissioned blockchain can maintain liveness under such powerful attacks. As shown in §7.3, existing consensus protocols, which ran on static committee nodes, lost liveness *until* the DoS attack ended. In contrast, each time after we attacked 100 nodes, EGES’s throughput had a temporary drop and recovered *before* the DoS attack ended, which shows that EGES can ensure practical liveness under such powerful attacks.

After the first attack, the line started to go up after 11.3s, much slower than the other attacks (about 3.1s). We inspected the log and found that the slow recovery was because the proposer for the next block happened to be attacked, and EGES waited until D more blocks to confirm that block as empty. After the second attack, the line took about 7.2s to go up. This is because most real acceptors happened to be attacked together with the proposer, which makes the arbiters failed to finalize the block for the proposer (§4.4). For the other three attacks, the arbiters successfully helped corresponding proposers to finalize their blocks quickly.

To evaluate EGES performance on network partitions, we manually divided the network into two partitions at 200s and reconnected them at 400s, with one partition containing 80% nodes and the other containing 20% nodes. Figure 6b shows the throughput measured in the large partition. Overall, the large partition maintained liveness during the partition. The small partition

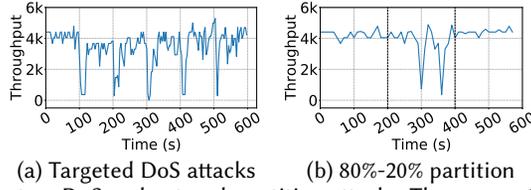


Fig. 6. EGES’s throughput on DoS and network partition attacks. There were 1000 nodes on AWS at 0s.

did not succeed in confirming any block during the partition and caught up after the network reconnected, preserving safety. There are two obvious throughput drops in the figure, which are caused by the pre-designated proposers being in the small partition, and EGES confirmed empty blocks for them. Note that EGES may temporarily lose liveness in catastrophic partitions (e.g., 50-50 or 40-30-30 partitions) but can preserve safety. §7.3 shows a quantitative analysis of how EGES can preserve liveness under network partitions.

7.3 Sensitivity

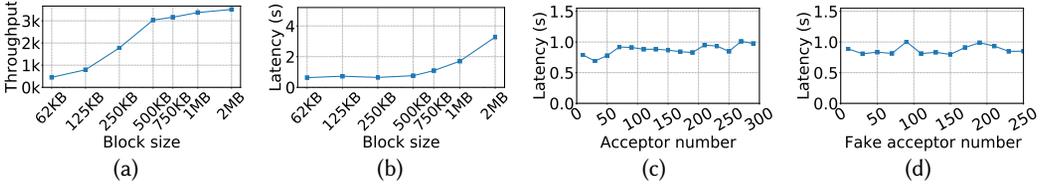


Fig. 7. Sensitivity on block size, acceptor numbers, and expected fake acceptor numbers (cluster setting).

Figure 7a and Figure 7b show EGES’s performance sensitivity on block size. When the block size was larger, EGES’s throughput did increase, but its block confirm latency also increased. In our evaluation, we set EGES’s block size to be 750KB, which is a near-optimal setting for both throughput and latency.

EGES’s throughput and confirm latency depend on three important protocol parameters, the number of acceptors, the number of fake acceptors, and block size. Figure 7c and Figure 7d show the sensitivity results. EGES’s performance turns out to be insensitive to the first two parameters because the latency is dominated by the time for broadcasting the new block on the P2P network.

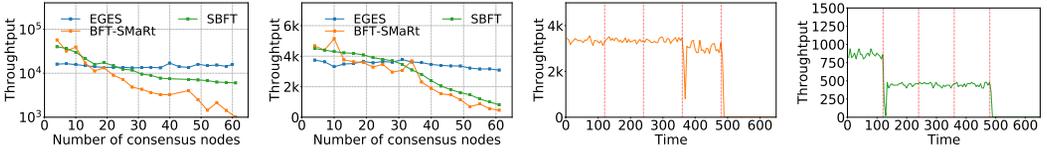
7.4 Comparison to BFT-SMaRt and SBFT

Since BFT-SMaRt with 10 (committee) nodes was faster than EGES with 100 acceptors, we evaluated both of them on a different number of nodes because more such nodes can tolerate more faults and DoS attacks. Figure 8a shows the results using the same setting for both systems (e.g., in our cluster, TC disabled, and the same number of transactions in each batch). Overall, EGES throughput was stable because the number of acceptors affects little on the latency in the seeking for quorum ACKs phase. BFT-SMaRt’s throughput drops dramatically because its protocol involves a quadratic number of messages on the number of ordering nodes.

Figure 8a also shows SBFT’s performance. In the non-geo-replicated mode, when the number of nodes increased from 4 to 62, SBFT’s throughput dropped from 38.2K to 6.9K transactions/s. This is because SBFT’s collectors (§2.2) need to collect more messages and to verify their signatures, so the time spent in collectors increased from 2.5ms to 13.1ms.

Figure 8b shows the performance comparison of EGES, BFT-SMaRt, and SBFT in the geo-replicated setting. EGES’s throughput was at least 3.4X larger than both systems on 62 nodes. BFT-SMaRt’s performance trend was similar to the no-delay setting because of its PBFT all-to-all broadcasted messages. SBFT’s throughput also dropped dramatically because some nodes became stragglers for the collectors due to the varied RTT. Since SBFT’s fast path can only tolerate a small number of

883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931



(a) No RTT delay ($t_c = 0$) (b) Geo-replicated mode. (c) BFT-SMaRt with 10 nodes (d) SBFT with 62 nodes
Fig. 8. Comparing EGES, SBFT, and BFT-SMaRt.

straggler nodes (usually two (§2.2)), we observed that 87% of the consensus rounds in SBFT have reverted to the slow path (PBFT).

More importantly, EGES can safely switch its acceptor groups across blocks, and it tolerated various failure scenarios, including DoS attacks (Figure 6a). For comparison, we evaluated the performance of BFT-SMaRt and SBFT on node failures (i.e., DoS attacks targeting consensus nodes). Figure 8c shows the result of BFT-SMaRt with its default 10-node setting. We randomly killed one node on each vertical line. The third time we killed its leader coincidentally, so there was a noticeable performance drop. BFT-SMaRt’s throughput dropped to zero after we killed the fourth node. For SBFT (Figure 8d), we started with 62 nodes and killed 7 nodes every time. Since SBFT’s fast path can only tolerate two crashed or straggler nodes, its throughput dropped significantly (reverted to PBFT) after the first kill.

Overall, EGES is complementary to BFT-SMaRt and SBFT: BFT-SMaRt is the fastest in a small scale; SBFT has better scalability, but its high performance requires a synchronous network (stated in their paper). EGES achieved reasonable efficiency and DoS resiliency in a geo-replicated setting.

7.5 Discussions

EGES has two limitations. First, EGES requires each node to have an SGX device. We deem this requirement reasonable because SGX is available on commodity hardware, and both academia and industry are actively improving the security of SGX. Recent permissionless [87] and permissioned blockchains [6, 66, 82] also use SGX. Second, EGES targets Internet-scale permissioned blockchain systems (e.g., a global payment system [16]), while for small-scale deployments (e.g., supply chain among a few small companies), existing consensus protocols (e.g., BFT-SMaRt) are more suitable.

Our paper reveals that, in addition to safety and high performance, DoS resistance is also an essential evaluation metric for a practical Internet-scale blockchain application, including e-voting [77], decentralized auction [19], and payment systems [16]. Moreover, the attested SGX enclave on each EGES node brings the potential to port existing centralized SGX-powered applications [51, 52, 67, 72] onto EGES and to make them DoS-resistant. For instance, ToR [30] is a popular anonymous network and is widely used for providing client anonymity to blockchain systems [33, 44]; SGX-ToR [51] greatly improves the security and privacy of ToR by leveraging SGX. However, SGX-ToR relies on a few directory servers for maintaining the list of attested nodes (relays), which has been shown [73] to be susceptible to DoS attacks. By deploying SGX-ToR’s directory service as a blockchain application on EGES, SGX-ToR can be made DoS-resistant.

8 CONCLUSION

We have presented EGES, the first efficient permissioned blockchain consensus protocol that can tolerate targeted DoS and partition attacks. EGES achieves comparable performance to existing fastest permissioned blockchain’s consensus protocols while achieving much stronger robustness. Our evaluation reveals that, in addition to safety and performance, DoS resistance should also be an essential evaluation metric for a blockchain system deployed on the Internet. EGES’s source code is available on github.com/performance21-p216/eges.

REFERENCES

- [1] [n.d.]. Algorand/go-algorand. <https://github.com/algorand/go-algorand/releases/tag/v2.0.14-beta>.
- [2] [n.d.]. hyperledger-labs/minbft. <https://github.com/hyperledger-labs/minbft>.
- [3] [n.d.]. OSI model - Wikipedia. https://en.wikipedia.org/wiki/OSI_model.
- [4] 2017. Blockchain for Supply Chain. <https://www.ibm.com/blockchain/industries/supply-chain>.
- [5] 2017. Medical Chain. <http://www.medicalchain.org/>.
- [6] 2019. *CCF: A Framework for Building Confidential Verifiable Replicated Services*. Technical Report MSR-TR-2019-16. Microsoft.
- [7] 2019. Cryptocurrency statistics.
- [8] 2021. web3.js - Ethereum JavaScript API. <https://web3js.readthedocs.io/>.
- [9] Mansoor Ahmed and Kari Kostiaainen. 2018. Identity Aging: Efficient Blockchain Consensus. *arXiv preprint arXiv:1804.07391* (2018).
- [10] Yackolley Amoussou-Guenou, Antonella Del Pozzo, Maria Potop-Butucaru, and Sara Tucci-Piergiorganni. 2018. Correctness of tendermint-core blockchains. In *22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, et al. 2018. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference (EuroSys 2018)*. ACM, 30.
- [12] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. 2017. Hijacking bitcoin: Routing attacks on cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 375–392.
- [13] Lăcrămioara Aștefanoaei, Pierre Chambart, Antonella Del Pozzo, Edward Tate, Sara Tucci, and Eugen Zălinescu. 2020. Tenderbake—Classical BFT Style Consensus for Public Blockchains. *arXiv preprint arXiv:2001.11965* (2020).
- [14] J Aumasson and L Merino. 2016. SGX Secure Enclaves in Practice—Security and Crypto Review. *Black Hat* (2016).
- [15] Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. 2018. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 913–930.
- [16] Mathieu Baudet, Avery Ching, Andrey Chursin, George Danezis, François Garillot, Zekun Li, Dahlia Malkhi, Oded Naor, Dmitri Perelman, and Alberto Sonnino. 2019. State machine replication in the Libra Blockchain.
- [17] Iddo Bentov, Yan Ji, Fan Zhang, Yunqi Li, Xueyuan Zhao, Lorenz Breidenbach, Philip Daian, and Ari Juels. 2017. Tesseract: Real-Time Cryptocurrency Exchange using Trusted Hardware. Cryptology ePrint Archive, Report 2017/1153. Accessed:2017-12-04.
- [18] Iddo Bentov, Rafael Pass, and Elaine Shi. 2016. Snow White: Provably Secure Proofs of Stake. <https://eprint.iacr.org/2016/919.pdf>. Accessed: 2016-11-08.
- [19] Erik-Oliver Blass and Florian Kerschbaum. 2017. Strain: A Secure Auction for Blockchains. Cryptology ePrint Archive, Report 2017/1044. Accessed:2017-11-06.
- [20] Ethan Buchman. 2016. Tendermint: Byzantine Fault Tolerance in the Age of Blockchains. http://atrium.lib.uoguelph.ca/xmlui/bitstream/handle/10214/9769/Buchman_Ethan_201606_MAsc.pdf. Accessed: 2017-02-06.
- [21] Vitalik Buterin. 2014. Ethereum: A next-generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper>. Accessed: 2016-08-22.
- [22] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI '99)*.
- [23] Raymond Cheng, Fan Zhang, Jernej Kos, Warren He, Nicholas Hynes, Noah Johnson, Ari Juels, Andrew Miller, and Dawn Song. 2018. Ekiden: A Platform for Confidentiality-Preserving, Trustworthy, and Performant Smart Contract Execution. *arXiv preprint arXiv:1804.05141* (2018).
- [24] Jin-Hee Cho, Dilli P Sharma, Hooman Alavizadeh, Seunghyun Yoon, Noam Ben-Asher, Terrence J Moore, Dong Seong Kim, Hyuk Lim, and Frederica F Nelson. 2020. Toward proactive, adaptive defense: A survey on moving target defense. *IEEE Communications Surveys & Tutorials* 22, 1 (2020), 709–745.
- [25] Byung-Gon Chun, Petros Maniatis, Scott Shenker, and John Kubiatowicz. 2007. Attested append-only memory: Making adversaries stick to their word. In *ACM SIGOPS Operating Systems Review*, Vol. 41. ACM, 189–204.
- [26] Victor Costan and Srinivas Devadas. 2016. Intel SGX Explained. *IACR Cryptology ePrint Archive* 2016 (2016), 86.
- [27] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2017. Ouroboros Praos: An adaptively-secure, semi-synchronous proof-of-stake protocol. Cryptology ePrint Archive, Report 2017/573. Accessed: 2017-06-29.
- [28] Christian Decker, Jochen Seidel, and Roger Wattenhofer. 2016. Bitcoin meets strong consistency. In *Proceedings of the 17th International Conference on Distributed Computing and Networking*. ACM, 13.
- [29] Deloitte. 2018. Blockchain Performance Report. https://www2.deloitte.com/content/dam/Deloitte/ie/Documents/Technology/IE_C_blockchain_performance_report.pdf.

- 1030 [30] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. *Tor: The second-generation onion router*. Technical Report.
1031 Naval Research Lab Washington DC.
- 1032 [31] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. 1988. Consensus in the presence of partial synchrony. *Journal of*
1033 *the ACM (JACM)* 35, 2, 288–323.
- 1034 [32] Ittay Eyal, Adem Efe Gencer, Emin Gun Sirer, and Robbert van Renesse. 2016. Bitcoin-NG: A Scalable Blockchain
1035 Association.
- 1036 [33] Qi Feng, Debiao He, Sherali Zeadally, Muhammad Khurram Khan, and Neeraj Kumar. 2019. A survey on privacy
1037 protection in blockchain system. *Journal of Network and Computer Applications* 126 (2019), 45–58.
- 1038 [34] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. 1985. Impossibility of distributed consensus with one faulty
1039 process. *Journal of the ACM (JACM)* 32, 2, 374–382.
- 1040 [35] Arthur Gervais, Ghassan O Karame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. 2016. On the
1041 Security and Performance of Proof of Work Blockchains. <https://eprint.iacr.org/2016/555.pdf>. Accessed: 2016-08-10.
- 1042 [36] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017. Algorand: Scaling Byzantine
1043 Agreements for Cryptocurrencies. Cryptology ePrint Archive, Report 2017/454. Accessed: 2017-06-29.
- 1044 [37] Seth Gilbert and Nancy Lynch. 2002. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant
1045 web services. *Acm Sigact News* 33, 2, 51–59.
- 1046 [38] Jinyu Gu, Zhichao Hua, Yubin Xia, Haibo Chen, Binyu Zang, Haibing Guan, and Jinming Li. 2017. Secure live migration
1047 of SGX enclaves on untrusted cloud. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and*
1048 *Networks (DSN)*. IEEE, 225–236.
- 1049 [39] Guy Golan Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael K Reiter, Dragos-Adrian
1050 Seredinschi, Orr Tamir, and Alin Tomescu. 2019. SBFT: a Scalable Decentralized Trust Infrastructure for Blockchains.
1051 IEEE/IFIP International Conference on Dependable System and Network (DSN 2019).
- 1052 [40] Indranil Gupta, A-M Kermarrec, and Ayalvadi J Ganesh. 2006. Efficient and adaptive epidemic-style protocols for reliable
1053 and scalable multicast. *IEEE Transactions on Parallel and Distributed Systems* 17, 7 (2006), 593–605.
- 1054 [41] Mike Hamburg, Paul Kocher, and Mark E Marson. 2012. Analysis of Intel’s Ivy Bridge digital random number generator.
1055 Online: http://www.cryptography.com/public/pdf/Intel_TRN_G_Report_20120312.pdf (2012).
- 1056 [42] Dominik Harz, Lewis Gudgeon, Arthur Gervais, and William J Knottenbelt. 2019. Balance: Dynamic adjustment of
1057 cryptocurrency deposits. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*.
1058 1485–1502.
- 1059 [43] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. 2015. Eclipse Attacks on Bitcoin’s Peer-to-Peer
1060 Network. In *24th USENIX Security Symposium (USENIX Security 15)*. 129–144.
- 1061 [44] Ryan Henry, Amir Herzberg, and Aniket Kate. 2018. Blockchain access privacy: Challenges and directions. *IEEE Security*
1062 *& Privacy* 16, 4 (2018), 38–45.
- 1063 [45] Charlie Hou, Mingxun Zhou, Yan Ji, Phil Daian, Florian Tramer, Giulia Fanti, and Ari Juels. 2019. SquirRL: Au-
1064 tomating Attack Analysis on Blockchain Incentive Mechanisms with Deep Reinforcement Learning. *arXiv preprint*
1065 *arXiv:1912.01798* (2019).
- 1066 [46] Jianyu Jiang, Xusheng Chen, Tsz-On Li, Cheng Wang, Tianxiang Shen, Shixiong Zhao, Heming Cui,
1067 Cho-Li Wang, and Fengwei Zhang. 2020. Uranus: Simple, Efficient SGX Programming and Its Applications. In
1068 *Proceedings of the 15th ACM on Asia Conference on Computer and Communications Security 2020 (ASIACCS ’20, accepted)*.
1069 <https://hemingcui.github.io/accepted/asiaccs20-uranus.pdf>.
- 1070 [47] Philipp Jovanovic. 2016. ByzCoin: Securely Scaling Blockchains. <http://hackingdistributed.com/2016/08/04/byzcoin/>.
1071 Accessed: 2019-08-01.
- 1072 [48] A-M Kermarrec, Laurent Massoulié, and Ayalvadi J. Ganesh. 2003. Probabilistic reliable dissemination in large-scale
1073 systems. *IEEE Transactions on Parallel and Distributed systems* 14, 3 (2003), 248–258.
- 1074 [49] Anne-Marie Kermarrec and Maarten Van Steen. 2007. Gossiping in distributed systems. *ACM SIGOPS Operating Systems*
1075 *Review* 41, 5 (2007), 2–7.
- 1076 [50] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2016. Ouroboros: A Provably Secure Proof-
1077 of-Stake Blockchain Protocol. <https://pdfs.semanticscholar.org/1c14/549f7ba7d6a000d79a7d12255eb11113e6fa.pdf>.
1078 Accessed: 2017-02-20.
- 1079 [51] Seong Min Kim, Juhyeng Han, Jaehyeong Ha, Taesoo Kim, and Dongsu Han. 2017. Enhancing Security and Privacy of
1080 Tor’s Ecosystem by Using Trusted Execution Environments.. In *NSDI*. 145–161.
- 1081 [52] Taehoon Kim, Joongun Park, Jaewook Woo, Seungheun Jeon, and Jaehyuk Huh. 2019. ShieldStore: Shielded In-memory
1082 Key-value Storage with SGX. In *Proceedings of the Fourteenth EuroSys Conference 2019*. ACM, 14.
- 1083 [53] Eleftherios Kokoris Kogias, Philipp Jovanovic, Nicolas Gailly, Ismail Khoffi, Linus Gasser, and Bryan Ford. 2016.
1084 Enhancing Bitcoin Security and Performance with Strong Consistency via Collective Signing. In *25th USENIX Security*
1085 *Symposium (USENIX Security 16)*. USENIX Association, Austin, TX.

- 1079 [54] Oleg Kupreev, Ekaterina Badovskaya, and Alexander Gutnikov. 2019. DDoS attacks in Q2 2019.
- 1080 [55] Leslie Lamport. [n.d.]. Paxos made simple. <http://research.microsoft.com/en-us/um/people/lamport/pubs/paxos-simple.pdf>.
- 1081 [56] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 4, 3, 382–401.
- 1082 [57] Joshua Lind, Oded Naor, Ittay Eyal, Florian Kelbert, Emin Gun Sirer, and Peter Pietzuch. 2019. Teechain: A Secure Payment Network with Asynchronous Blockchain Access. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles* (Huntsville, Ontario, Canada) (SOSP '19). ACM, New York, NY, USA, 63–79. <https://doi.org/10.1145/3341301.3359627>
- 1083 [58] Joshua Lind, Christian Priebe, Divya Muthukumar, Dan O’Keeffe, Pierre-Louis Aublin, Florian Kelbert, Tobias Reiher, David Goltzsche, David Eysers, Rüdiger Kapitza, et al. 2017. Glamdring: Automatic application partitioning for Intel SGX. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, Santa Clara, CA.
- 1084 [59] Sinisa Matetic, Mansoor Ahmed, Kari Kostiaainen, Aritra Dhar, David Sommer, Arthur Gervais, Ari Juels, and Srdjan Capkun. 2017. ROTE: Rollback Protection for Trusted Execution. *IACR Cryptology ePrint Archive 2017* (2017), 48.
- 1085 [60] David Mazieres. 2007. *Paxos made practical*. Technical Report. Technical report, 2007. <http://www.scs.stanford.edu/dm/home/papers>.
- 1086 [61] Gabriel Mendonça, Gustavo HA Santos, Edmundo de Souza e Silva, Rosa MM Leão, Daniel S Menasché, and Don Towsley. 2019. An extremely lightweight approach for ddos detection at home gateways. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 5012–5021.
- 1087 [62] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The honey badger of BFT protocols. <https://eprint.iacr.org/2016/199.pdf>. Accessed: 2017-01-10.
- 1088 [63] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. <https://bitcoin.org/bitcoin.pdf>. Accessed: 2015-07-01.
- 1089 [64] Christopher Natoli and Vincent Gramoli. 2016. The Balance Attack Against Proof-Of-Work Blockchains: The R3 Testbed as an Example. <http://arxiv.org/abs/1612.09426>. Accessed: 2017-02-15.
- 1090 [65] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proceedings of the USENIX Annual Technical Conference (ATC '14)*.
- 1091 [66] poet [n.d.]. <https://www.hyperledger.org/projects/sawtooth>.
- 1092 [67] Christian Priebe, Kapil Vaswani, and Manuel Costa. 2018. EnclaveDB: A Secure Database using SGX. In *Proceedings of the 2018 IEEE Symposium on Security and Privacy*. IEEE, 0.
- 1093 [68] Seon Pyo Kim. 2020. Tenderand: Randomized leader election in Tendermint. (2020).
- 1094 [69] Alison Rayome. 2018. Major DDoS attack lasts 297 hours, as botnets bombard businesses.
- 1095 [70] Robert Riemann and Stéphane Grumbach. 2017. Distributed Protocols at the Rescue for Trustworthy Online Voting. arXiv:1705.04480. <https://doi.org/10.5220/0006228504990505> Accessed: 2017-06-29.
- 1096 [71] Devavrat Shah et al. 2009. Foundations and Trends® in Networking. *Foundations and Trends® in Networking* 3, 1 (2009), 1–125.
- 1097 [72] Fahad Shaon, Murat Kantarcioglu, Zhiqiang Lin, and Latifur Khan. 2017. SGX-BigMatrix: A Practical Encrypted Data Analytic Framework With Trusted Processors. In *Proceedings of the 17th ACM conference on Computer and communications security (CCS '10)*.
- 1098 [73] Tianxiang Shen, Jianyu Jiang, Yunpeng Jiang, Xusheng Chen, Ji Qi, Shixiong Zhao, Fengwei Zhang, Xiapu Luo, and Heming Cui. 2021. DAENet: Making Strong Anonymity Scale in a Fully Decentralized Network. *IEEE Transactions on Dependable and Secure Computing* (2021).
- 1099 [74] Yonatan Sompolinsky and Aviv Zohar. 2013. Accelerating Bitcoin’s Transaction Processing. Fast Money Grows on Trees, Not Chains. , 881 pages. <http://eprint.iacr.org/2013/881.pdf>
- 1100 [75] João Sousa, Alysson Bessani, and Marko Vukolić. 2017. A Byzantine Fault-Tolerant Ordering Service for the Hyperledger Fabric Blockchain Platform. IEEE/IFIP International Conference on Dependable System and Network (DSN 2018).
- 1101 [76] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. 2001. Chord: A scalable peer-to-peer lookup service for internet applications. *SIGCOMM Comput. Commun. Rev.* 31, 4 (2001), 149–160.
- 1102 [77] Pavel Tarasov and Hitesh Tewari. 2017. Internet Voting Using Zcash. Cryptology ePrint Archive, Report 2017/585. Accessed: 2017-06-29.
- 1103 [78] Sergei Tikhomirov. 2017. Ethereum: state of knowledge and research perspectives. Accessed:2018-01-05.
- 1104 [79] Saar Tochner, Aviv Zohar, and Stefan Schmid. 2020. Route Hijacking and DoS in Off-Chain Networks. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 228–240.
- 1105 [80] Muoi Tran, Loi Luu, Min Suk Kang, Iddo Bentov, and Prateek Saxena. 2017. Obscuro: A Bitcoin Mixer using Trusted Execution Environments. Cryptology ePrint Archive, Report 2017/974. Accessed:2017-10-06.
- 1106 [81] Sridhar Venkatesan, Massimiliano Albanese, Kareem Amin, Sushil Jajodia, and Mason Wright. 2016. A moving target defense approach to mitigate DDoS attacks against proxy-based architectures. In *2016 IEEE conference on*
- 1107
- 1108
- 1109
- 1110
- 1111
- 1112
- 1113
- 1114
- 1115
- 1116
- 1117
- 1118
- 1119
- 1120
- 1121
- 1122
- 1123
- 1124
- 1125
- 1126
- 1127

- 1128 *communications and network security (CNS)*. IEEE, 198–206.
- 1129 [82] Giuliana Santos Veronese, Miguel Correia, Alysson Neves Bessani, Lau Cheuk Lung, and Paulo Verissimo. 2013. Efficient
1130 byzantine fault-tolerance. *IEEE Trans. Comput.* 62, 1, 16–30.
- 1131 [83] TY Wong, KT Law, John CS Lui, and Man Hon Wong. 2006. An efficient distributed algorithm to identify and traceback
1132 ddos traffic. *Comput. J.* 49, 4 (2006), 418–442.
- 1133 [84] Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. 2019. HotStuff: BFT consensus with
1134 linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM,
1135 347–356.
- 1136 [85] Rui Yuan, Yu-Bin Xia, Hai-Bo Chen, Bin-Yu Zang, and Jan Xie. 2018. ShadowEth: Private Smart Contract on Public
1137 Blockchain. *Journal of Computer Science and Technology* 33, 3 (2018), 542–556.
- 1138 [86] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. 2016. Town Crier: An Authenticated Data Feed
1139 for Smart Contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM,
1140 270–282.
- 1141 [87] Fan Zhang, Ittay Eyal, Robert Escriva, Ari Juels, and Robbert van Renesse. 2017. REM: Resource-Efficient Mining for
1142 Blockchains. <http://eprint.iacr.org/2017/179>.
- 1143 [88] Qichao Zhang, Zhuyun Qi, Xiaoyou Liu, Tao Sun, and Kai Lei. 2018. Research and Application of BFT Algorithms Based
1144 on the Hybrid Fault Model. In *2018 1st IEEE International Conference on Hot Information-Centric Networking (HotICN)*.
1145 IEEE, 114–120.

1145 Appendices

1146 A PROOF OF SAFETY

1147 EGES ensures safety with overwhelming probability (i.e., $> 1 - 10^{-10}$). Formally, if a node confirms a
1148 block b as the i^{th} block on the blockchain, the probability that another member node confirms $b' \neq b$
1149 as the i^{th} block is $< 10^{-10}$.

1150 We prove the safety guarantee of EGES by induction: suppose EGES guarantees safety from the 0^{th}
1151 block to the $(n-1)^{\text{th}}$ block ([hypothesis 1](#)), and we prove that there is only one unique block that can
1152 be confirmed as the n^{th} block among nodes in the blockchain. The base case is trivial because all
1153 nodes start from the same 0^{th} block.

1154 **LEMMA A.1.** *if two nodes have the same maximum confirmed block in their chain (i.e., $MC = n-1$ due
1155 to [hypothesis 1](#)), then during consensus for the $(n+i)^{\text{th}}$ block where $i \geq 0$, as long as MC is not changed,
1156 these two nodes see the same member list.*

1157 **PROOF.** Proving this lemma is trivial if EGES works on a fixed member list, and we will show in
1158 [Appendix D](#) that EGES’s protocol for dynamic memberships also ensures this lemma. \square

1159 **LEMMA A.2.** (*Invariant 1 in §4.1*): *at most one proposal can be generated for the n^{th} block.*

1160 **PROOF.** This lemma is proved by two steps. First, as the proposer for the n^{th} block is encrypted in
1161 the $(n-lb)^{\text{th}}$ block, and the $(n-lb)^{\text{th}}$ block is the same among nodes because of hypothesis 1.
1162 Therefore, there is only one proposer (may have failed) for the n^{th} block. Second, this proposer
1163 generates at most one proposal, and non-proposer nodes cannot generate valid proposals for the n^{th}
1164 block because EGES’s consensus module runs in SGX. \square

1165 **Proof of the induction step.** In EGES, each block has only two choices ([Lemma 2](#)), and a confirmed
1166 block must be first finalized ([§4.1](#)). Therefore, it is sufficient to prove the following [proposition 1](#): the
1167 probability that one node finalizes empty_n (event A), and another node finalizes proposal_n (event B)
1168 is overwhelmingly small.

1169 For event A, suppose node X finalizes empty_n . We use f_{m_x} to denote the maximum finalized
1170 block index on node X . Consider blocks with indices in $[n+1, f_{m_x}]$. Since EGES finalizes *empty*
1171

1177 blocks in descending order (Algorithm 2 Line 12~15), there are no undecided blocks in $[n + 1, f_{m_x}]$,
 1178 and we can have another level of induction by supposing blocks finalized as empty in $(n + 1, f_{m_x}]$
 1179 are finalized consistently (name it hypothesis 2). For event B, proposal_n can be finalized either
 1180 by P_n (call it event B1) or by subsequent proposers that have learnt this proposal (event B2).

1181 First, we prove that the probability that event A and event B1 happen together is overwhelmingly
 1182 low. Suppose that a portion p of all M EGES nodes received and cached the proposal_n , and we
 1183 calculate the probability for event B1. We use R_e to denote the number of acceptors for the n^{th} block
 1184 that RReceived proposal_n . Since proposal_n is broadcasted in EGES's P2P network and the stealth
 1185 acceptors are selected uniformly, R_e follows hypergeometric distribution $R_e \sim H(M, n_A, p \times M)$.
 1186 Thus, the probability that P_n finalizes proposal_n is

$$1187 \quad \text{Prob}(B1) = \text{Prob}(R_e > \tau \times n_A)$$

1188 We then calculate the probability of event A. Event A infers that after the n^{th} block, there are at least
 1189 D non-empty blocks that are finalized and carrying n in the undecided list. This means each proposer
 1190 of these D blocks received $(\tau \times n_A)$ ACKs from their acceptor group, and none of these acceptors
 1191 sending those ACKs has received proposal_n . For each of the D blocks, the number of acceptors
 1192 NR not receiving proposal_n follows hypergeometric distribution $NR \sim H(M, n_A, (1 - p) \times M)$.
 1193 Therefore, the probability of event A is

$$1194 \quad \text{Prob}(A) = (\text{Prob}(NR > \tau \times n_A))^D$$

1195 The calculation shows that the probability of event A and event B1 happening together $\text{Prob}(A) \times$
 1196 $\text{Prob}(B1)$ is overwhelmingly low for any delivery rate p by setting τ and n_A (§7.3). For instance, our
 1197 evaluation chose τ as 59%, D as 4, n_A as 300, M as 10K, and the probability of EGES enforcing safety is
 1198 $1 - 10^{-9}$. In real deployments, M may change due to membership changes; however, when M is much
 1199 bigger (e.g., 20X) than n_A , this probability is not sensitive to M because hypergeometric distribution
 1200 is approximate to binomial.

1201 For the second step, we prove that event A and event B2 cannot happen together. For event
 1202 B2, we suppose that proposer P_i , where $i > n$, learns and finalizes proposal_n . We discuss by
 1203 comparing i and f_{m_x} and derive contradictions. If $i \leq f_{m_x}$, hypothesis 2 infers that X did not
 1204 finalize empty_i , so proposal_n is finalized together with proposal_i at node X , contradicting
 1205 to event A. Else if $i > f_{m_x}$, since a proposer can only finalize the maximum index in its local
 1206 \cup list (Algorithm 1 Line 9), for node P_i we can predicate that blocks with index in $[n + 1, i)$
 1207 are finalized. Due to hypothesis 2, blocks within $[n + 1, f_{m_x}]$ are finalized the same as node X ,
 1208 and therefore P_i should also finalize the n^{th} block as empty, causing contradiction.

1209 Putting the two steps together, we proved proposition 1 and thus proved the induction step that
 1210 the n^{th} block must be confirmed consistently among nodes with overwhelmingly high probability.
 1211 Therefore, EGES ensures safety with overwhelmingly high probability.
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225

B PARAMETER SELECTION

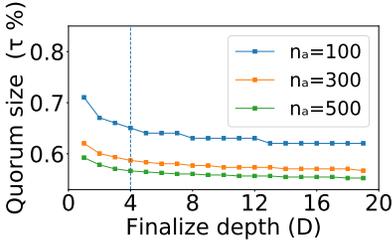


Fig. 9. Parameter selection for τ and D on 10K nodes for different n_A values.

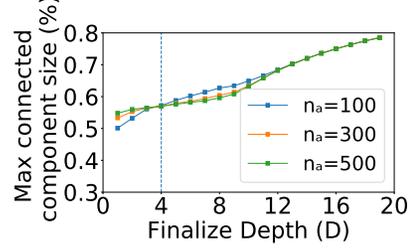


Fig. 10. Connected component size required to ensure liveness with different D values.

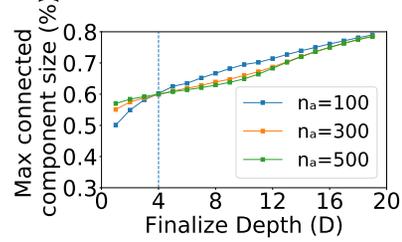
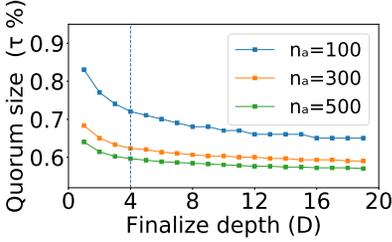


Fig. 11. Parameter selection and liveness requirements if EGES lets each node to independently decide its committee membership.

Figure 9 shows the relation between τ and D to ensure EGES’s safety. With a smaller τ , a proposer P_n can finalize proposal _{n} after collecting fewer ACKs from acceptors, so subsequent proposers need more rounds of checking (larger D) when trying to finalize the empty _{n} (§4.3).

The τ and D values also affect EGES’s ability to achieve liveness (confirm non-empty blocks) on network partitions (or ubiquitous DoS attacks). We quantify this ability to the “minimum largest connected component size” ($cc\%$) required in the P2P graph, provided that nodes in a connected component can reach each other before a timeout. A smaller cc means that EGES is more robust to partition and ubiquitous DoS attacks. From a mathematical aspect, as long as the probability of finalizing a proposal is non-zero, the probability p_D that D consecutive proposals are successfully finalized is always larger than zero, inferring that *eventually* EGES can achieve liveness. However, we conservatively calculate the required cc to make the p_D larger than 5% for practical liveness, as shown in Figure 10. In our evaluation, we chose τ as 59%, D as 4, n_A as 300, which ensures both safety and achieves good liveness on partition attacks (§7.2).

Figure 11 shows the parameter selections if EGES does not use its stealth committee mechanism, but lets each node *independently* determine whether it is an acceptor with the probability of M/n_A , with n_A being the *expected* number of acceptors for each block. If EGES makes such a design choice, the Re (§5.1) becomes a binomial distribution with the probability of $p \times M/n_A$, and other distribution changes similarly. As shown in Figure 11, EGES would need a larger quorum size $\tau \times n_A$ and achieve worse liveness on network partition.

C EXAMPLES OF EGES'S CONSENSUS PROTOCOL

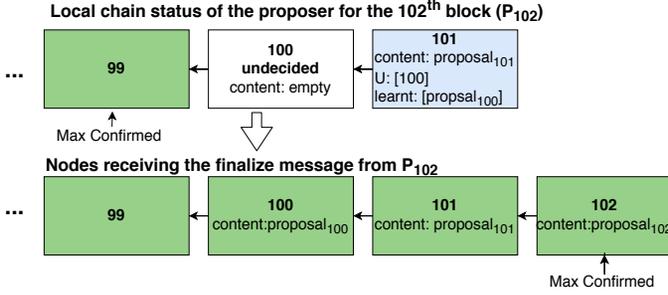


Fig. 12. An example where p_{102} helps finalizing proposal₁₀₀ while proposing its (102th) block.

Figure 12 gives an example illustrating how a proposer P_{102} helps to finalize an undecided proposal₁₀₀. Suppose P_{100} failed just before broadcasting its finalize message. Therefore, the 100th block's state is undecided among all nodes. Then, P_{101} learns proposal₁₀₀ and carries it in its finalize message, and P_{102} learns it. Then P_{102} finalizes proposal₁₀₀ together with proposal₁₀₂. Moreover, since all blocks before 102 are finalized, the chain is confirmed up to 102.

Figure 13 gives an example showing how an undecided block is finalized as empty. Suppose P_{200} failed before broadcasting its proposal, and $D = 4$. When $P_{201} \sim P_{204}$ asks whether their acceptors ($\mathbb{A}_{201} \sim \mathbb{A}_{204}$) receive proposal₂₀₀, they get no positive answers. Therefore, these four blocks are all finalized, carrying a message stating that four samplings have been conducted on the delivery rate of proposal₂₀₀, but no replied node has received proposal₂₀₀. This indicates that the probability that proposal₂₀₀ is finalized at some nodes is overwhelmingly low. Therefore, a node can independently finalize empty₂₀₀, after which the chain is confirmed to 204.

Figure 14 shows why it is essential that a checking mode proposer can only finalize proposal _{u_m} where $u_m = \max(\mathbb{U})$. Suppose we remove this restriction, and we consider the following scenario. (1) P_{200} fails after broadcasting proposal₂₀₀, and only very few nodes received it: none of $P_{201} \sim P_{204}$ learns proposal₂₀₀. (2) The network is divided into two partitions A&B just before P_{204} broadcasting its finalize message; P_{204} is in partition A, so nodes in partition A confirm proposal₂₀₄ and confirm empty₂₀₀. (3) P_{205} and P_{206} are in partition B, so they timeout waiting for the finalize message for proposal₂₀₄ and mark the 204th block as undecided. (4) P_{205} learns proposal₂₀₀ from one node from \mathbb{A}_{205} (who happens to be in the very few nodes) and carries proposal₂₀₀ with its finalize message, which is learnt by P_{206} . (5) P_{206} finalizes proposal₂₀₀ and causes inconsistency:

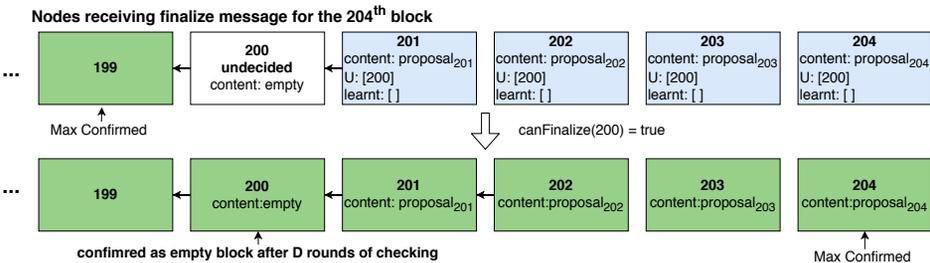
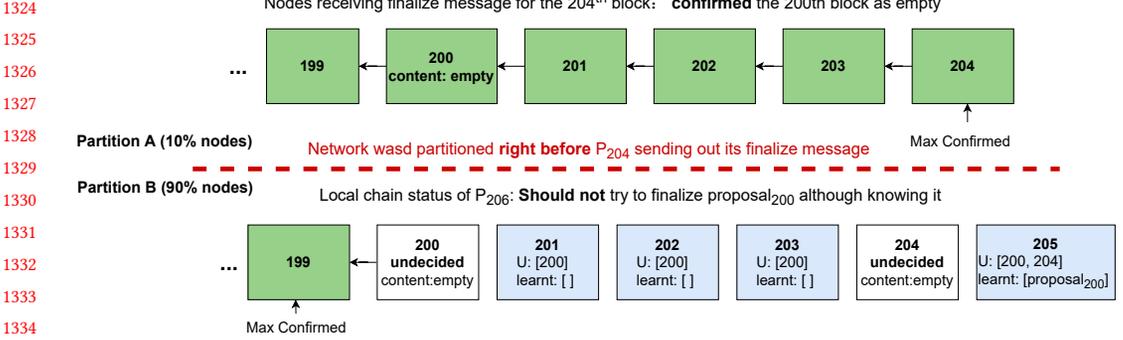


Fig. 13. An example for confirming an empty block (200th) after $D = 4$ succeeding blocks containing 200 in \mathbb{U} .



1336 Fig. 14. An example showing why a checking mode proposer can finalize only proposal_i where $i = \max(\mathbb{U})$.
 1337 $D = 4$ in this example.

1338 nodes in partition A confirm empty_{200} , while nodes in partition B confirm proposal_{200} .

1339 By restricting that a checking mode proposer can only finalize proposal_{u_m} where $u_m = \max(\mathbb{U})$,
 1340 such inconsistency will not happen. This is because nodes in partition B must first finalize the
 1341 empty_{204} before finalizing proposal_{200} . However, since proposal_{204} has already been delivered to
 1342 a large portion of nodes, this inconsistency cannot happen.

1343 D MEMBERSHIP AND ATTESTATION

1344 To support dynamic memberships, EGES leverages the idea from SCIFER [9] to record the joining of
 1345 new nodes as transactions on the blockchain. This mechanism ensures Lemma 1 because all updates
 1346 to the member list are only determined by confirmed blocks.

1347 When a node i wants to join the system, it needs to find a member node j to do attention (§2)
 1348 through an out-of-band peer discovery service. We assume that node i knows the genesis (0^{th}) block,
 1349 so node i can inductively verify the blockchain, without relying on whether peer j is malicious.

1350 A node i joins EGES with three steps. First, i launches its EGES enclave, which generates its account
 1351 (pk_i, sk_i) and creates the hardware monotonic counter c for defending forking attacks (Appendix E).
 1352 The node's account is securely saved to permanent storage using SGX's seal mechanism [26] for
 1353 recoveries from machine failures (e.g., power off). Then, i sends a *join* request to j . Second, j does
 1354 a standard SGX remote attestation [26], which succeeds with a signed quote Q_i from Intel's attestation
 1355 service, and i 's enclave transfers its public key pk_i and counter value c to j 's enclave through the
 1356 secure communication channel between two enclaves created during attestation. Third, node j 's
 1357 enclave creates a signed registration transaction including $pk_i, c, addr_i, Q_i$ and i 's ip address. Node i
 1358 joins EGES when the transaction is included in a confirmed block.

1359 E ENCLAVE INTERACTIONS

1360 Figure 15 shows the implementation of EGES enclave. An EGES enclave holds three data structures
 1361 shared among ECalls: `cache`, `is_proposer`, and `is_acceptor`, each as a hash map. As explained
 1362 in §4.1, the `cache` saves received block proposals. In our implementation, the `cache` only keeps
 1363 the hash values of block headers instead of whole blocks to save enclave memory. `is_proposer`
 1364 and `is_acceptor` are hash maps with block indices as keys and boolean as values, saving whether
 1365 this node will be in the committee for a future block.

1366 In Figure 15a, when a node's blockchain core module confirms the $(n - lb)^{\text{th}}$ block, it asyn-
 1367 chronously invokes an ECall letting the enclave check whether it will be the proposer/acceptor
 1368 for the n^{th} block (§4.2). In Figure 15b, when a node's core module appends the $(n - 1)^{\text{th}}$ block, it

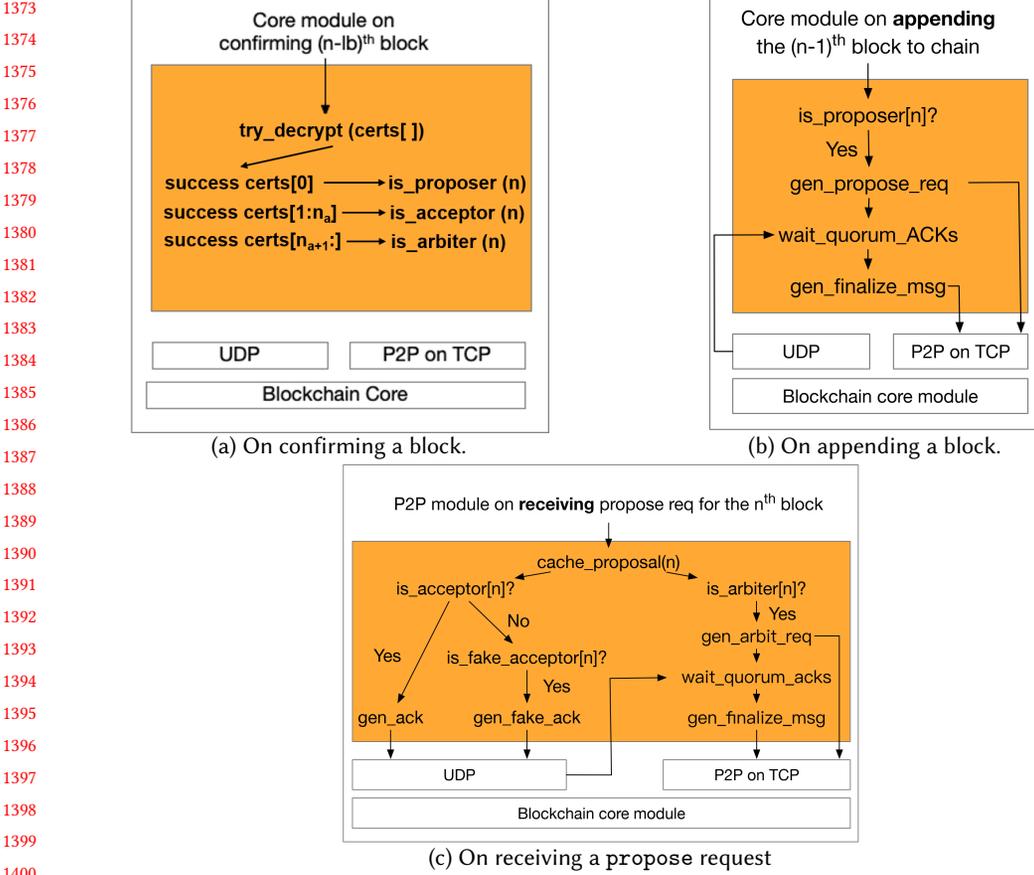


Fig. 15. EGES's enclave interactions (ECalls and OCalls). The enclave module is shaded in orange.

invokes an ECall with a batch of transactions, and the enclave will follow the protocol in Algorithm 1 if it is the proposer for the n^{th} block. In Figure 15c, when a node's P2P module received a propose request, it invokes an ECall passing this request to the enclave, and the enclave will generate an ACK that will be sent through UDP using an OCall if it is an acceptor (Algorithm 3) or fake acceptor (Algorithm 2 Line 4). If it is an arbiter for the n^{th} block, it will also start working as an arbiter (§4.4).

Enclave forking attacks. In the P2P scenario, one challenge is enclave forking attacks [59]. EGES must permit a node to reuse its sealed account (Appendix D) in case the node restarts its machine. However, a malicious node can create multiple copies of EGES enclaves with the same account pk , directs different messages to them, and lets them generate conflicting messages (e.g., block proposals). Existing defending techniques [38, 59] work in the client-server manner, where clients attest and communicate to only a single server. These techniques are not suitable for P2P settings because they will need every two EGES nodes to connect and attest each other.

EGES defends such attacks using SGX's platform counter [26], which is monotonic among all enclaves on the same machine. When a node launches its EGES enclave, the enclave increments and reads this counter value c and encloses this c to its registration transaction: the node's membership is bound to the enclave with counter value c but not the account pk . When the enclave sees a registration with the same account but a higher counter value, it quits automatically.