Efficient and DoS-resistant Consensus for Permissioned Blockchains

Xusheng Chen, Shixiong Zhao, Ji Qi, Jianyu Jiang, Haoze Song, Cheng Wang, Tsz On Li, T.-H. Hubert Chan, Fengwei Zhang, Xiapu Luo, Sen Wang, Gong Zhang, and Heming Cui.

Abstract—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 in 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 work proposes scalable BFT protocols that run on a larger number of committee nodes, their efficiency drops dramatically when only a small number of nodes are attacked.

In this paper, our EGES protocol leverages Intel SGX to develop a new abstraction called "stealth committee", which effectively hides the committee nodes into a large pool of fake committee nodes. EGES selects a distinct group of stealth committee for each block and confirms the same sequence of blocks among all nodes with overwhelming probability. Evaluation on typical geo-distributed settings shows that: (1) EGES is the first permissioned blockchain's consensus protocol that can tolerate tough DoS and network partition attacks; and (2) EGES achieves comparable throughput and latency as existing permissioned blockchains' protocols.

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 the same sequence of blocks (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 attached with a cryptocurrency mechanism (e.g., Bitcoin [1]) to incite nodes to follow the blockchain's protocol.

In contrast, a permissioned blockchain runs on a set of identified member nodes and can leverage the mature Byzantine Fault-Tolerant (BFT) protocols [2], [3], [4], [5] to achieve better efficiency (i.e., throughput and latency on confirming blocks). This paper focuses on permissioned blockchains because their decoupling from cryptocurrencies has facilitated the deployment of many real-world general data-sharing applications, including a medical chain among UK hospitals [6], IBM supply chains [7], and the Libra payment system [8].

For performance and regulation reasons (e.g., meeting the honesty threshold of BFT protocols [9]), a permissioned blockchain (e.g., Hyperledger Fabric [10]) 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 [7].

Unfortunately, as permissioned blockchains become popular and are deployed in large scales on the Internet, this static committee approach is vulnerable to targeted Denial-of-Service (DoS) and network partition [11], [12], [13] attacks. For instance, Libra [8] aims to build a global payment system (there are more than 5,000 banks in the US [14]) and identifies DoS attacks as a significant threat, but it only provides partial mitigation (§2.2).

Indeed, great progress has been made in designing scalable BFT protocols (e.g., SBFT [4]) 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 [15], [16], 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 [11], [17]. These systems select a distinct committee for each block to handle the frequent leaving of nodes and to provide fairness and can ensure liveness even if a committee fails to confirm a block.

Simply applying the dynamic committee approach, however, cannot ensure DoS-resistance. Another crucial requirement is unpredictability: the identifies of nodes in a committee must be unpredictable to the attacker before the committee tries to achieve consensus on a block. Otherwise, the attacker can adaptively attack the ready-to-be committee and cause the system stuck. For instance, ByzCoin [18] lets the proof-of-work winners of recent blocks be the committee, but these explicit nodes are easily targeted by a DoS attacker, causing ByzCoin to lose liveness permanently [19].

To the best of our knowledge, Algorand is the only work that can tolerate targeted DoS attacks. However, as Algorand is designed for permissionless blockchains, it confirms a block with up to 15 rounds and minute-level latency (discussed in §2.2), making it unsuitable for general data-sharing applications on permissioned blockchains (e.g., Libra).

This paper aims to explore the new design point of building a permissioned blockchain's consensus protocol that adopts the unpredictable dynamic committee merit to defend against targeted DoS or partition attacks, and at the same time, achieves comparable efficiency as existing BFT protocols (e.g., SBFT [4] has second-level latency).

A main obstacle is to ensure that any selected committee

meets the honesty requirement for byzantine problems: for consistency, each committee must have at most one-third of nodes being malicious [9]. Permissionless blockchains meet this requirement by selecting committees based on nodes' wealth in the built-in cryptocurrency mechanism (i.e., proof-of-stake), but cryptocurrencies are usually unavailable in permissioned blockchains. Consequently, to meet such a requirement, one has to unrealistically assume that almost all member nodes (>90%) are honest (calculated in §2).

Fortunately, recent work (e.g., Microsoft CCF [20], REM[21]) shows that the code integrity feature of SGX [22] can regulate the misbehavior of blockchain nodes. For instance, a recent implementation [23] of MinBFT leverages SGX to ensure that a node cannot send conflicting messages to different nodes and is incorporated into Hyperledger [10].

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

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

However, even equipped with SGX and stealth committee, it is still challenging to efficiently ensure both consistency (i.e., no two member nodes confirm conflicting blocks) and reasonable liveness (i.e., allow non-empty blocks to get confirmed) in the asynchronous Internet due to the FLP impossibility [24]. Specifically, suppose a committee node x for the n^{th} block fails to receive the $(n - 1)^{th}$ block after a timeout, x cannot distinguish whether it is because the committee for the $(n - 1)^{th}$ block failed to confirm the $(n - 1)^{th}$ block, or because x itself does not receive the committee nodes for the $(n - 1)^{th}$ 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 using 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 [11]. In the previous example, EGES confirms the proposal for the $(n - 1)^{th}$ block only if the proposal is delivered to a large portion of member nodes; if multiple rounds of the sampling show that very few nodes

protocol	DoS	with	consensus	number	tput	confirm
name	resistanc	SGX?	model	of nodes	(txn/s)	latency (s)
FCES	high	Yes	hybrid	300	3226	0.91
LGE5			nybiid	10k	2654	1.13
Algorand	high	No	BFT	10K	\sim 727	$\sim 22s$
PoET	high	Yes	Hybrid	100	149	45.2
Ethereum	high	No	BFT	100	178	82.3
SBFT	medium	No	BFT	62	1523	1.13
MinBFT	low	Yes	Hybrid	64	2478	0.80
BFT-SMaRt	low	No	BFT	10	4512	0.67
Tendermint	low	No	BFT	64	2462	1.31
HotStuff	low	No	BFT	64	2686	2.63
HoneyBadger	low	No	BFT	32	1078	9.39

TABLE 1: Comparison of EGES to baseline protocols. Analysis of DoS resistance is in §2; evaluation setup is covered in §8.

have received that proposal for the $(n-1)^{th}$ block, nodes in EGES consistently confirm the $(n-1)^{th}$ 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 [25], [26]. EGES unpredictably replaces the critical components (i.e., the committee) in a large system so that a DoS attacker cannot launch effective attacks targeting these components. EGES's consensus protocol also achieves good efficiency because confirming a block in a gracious run (e.g., the proposer can reach most acceptors) only involves two P2P broadcasts and a half UDP round trip (§5). We provide a rigorous analysis of EGES's DoS-resistance and proof of EGES's consistency guarantee in §6.

We implemented EGES using the codebase from Ethereum [27] and compared EGES with nine consensus protocols for blockchain systems, including five state-ofthe-art efficient BFT protocols for permissioned blockchains (BFT-SMaRt [28], SBFT [4], HoneyBadger [29], and HotStuff [5]), two SGX-powered consensus protocols for permissioned blockchains (Intel-PoET [30] and MinBFT [31]), the default consensus protocol in our codebase (Ethereum-PoW [27]), and two permissionless blockchains' protocols that run on dynamic committees (Algorand [11] and Tendermint [17]). We ran EGES on both our cluster and AWS. Evaluation shows 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 (§6) and evaluation (§8.2).
- 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.

^{1.} EGES stands for Efficient, GEneral, and Scalable consensus.

Compared to existing BFT protocols [28], [4], [29], [5] and SGX-powered consensus protocols [30], [31] 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 two-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. EGES's source code is available on github.com/hku-systems/eges. EGES can facilitate the deployments of various mission-critical, DoS-resistant permissioned blockchain applications on the Internet (e.g., evoting [32] and payment [8]).

In the rest of the paper, §2 introduces EGES's background and motivation; §3 defines the model of EGES; §4 gives a high-level overview of EGES; §5 introduces EGES's consensus protocol; §6 analyzes the safety and liveness of EGES; §8 shows our evaluation, and §9 concludes.

2 BACKGROUND AND RELATED WORK

We discuss targeted DoS and network partition attacks together in this paper because these two attacks cannot be effectively distinguished in an asynchronous network. Particularly, 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 these two nodes are partitioned in the network. Therefore, EGES maintains consistency by handling both cases together.

When discussing targeted DoS or partition attacks in this subsection, we assume that the attacker has an attack budget *B* (e.g., B = 300): the attacker can adaptively target *B* nodes at a time. This model is the same as Algorand's, which we will formally define in §3.

2.1 Intel SGX

Intel Software Guard eXtension (SGX) [22] is a hardware feature on commodity CPUs. SGX provides a secure execution environment called 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 [22] 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 [22]. Previous studies show that this random source complies with security and cryptographic standards and cannot be seen or tampered with from outside enclaves [33], [34].

Recent work shows that SGX can be leveraged to improve diverse aspects of blockchain systems. Intel's PoET [30] replaces the PoW puzzles with a trusted timer in SGX; EGES is more efficient than PoET (§8.1). REM [21] uses SGX to replace the useless PoW puzzles with useful computation (e.g., big data), orthogonal to EGES. Microsoft CCF [20] (originally named CoCo) is a permissioned blockchain platform using SGX to achieve transaction privacy, but it does not include a DoS-resistance approach. Scifer [35] uses SGX's attestation to establish reliable identities of nodes and maintains their identities on the blockchain, which is adopted in EGES (§7.1).

Ekiden [36] and ShadowEth [37] offload the execution of smart contracts to SGX-powered nodes to avoid redundant execution and to preserve privacy; TEEChain [38] uses SGX to build an efficient and secure off-chain payment channel; Town Crier [39] uses SGX to build a trustworthy data source for smart contracts; Tesseract [40] uses SGX to build a crosschain coin exchange framework; Obscuro [41] uses SGX to improve bitcoin's privacy; these systems do not focus on consensus protocols and are orthogonal to EGES.

2.2 Consensus for Permissioned Blockchains

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

BFT-SMaRt [2] is an optimized implementation of PBFT [3]. As each node broadcasts consensus messages to all other nodes, BFT-SMaRt has $O(n^2)$ message complexity to the number of committee nodes, resulting in poor scalability. Its paper [2] only evaluated up to 10 nodes. SBFT [4] is a scalable BFT protocol that uses a new type of committee nodes called collectors. A node sends its consensus messages to only c (usually c < 8) explicit collectors who will then broadcast a combined message using the threshold signature. SBFT's fast path can commit a block if fewer than c nodes have failed; however, SBFT's performance drops dramatically if an attacker targets the c collectors (§8.4).

HotStuff [5] is a BFT protocol optimized for frequent leader changes, and Libra [8] leverages Hotstuff to tolerate targeted DoS attacks *on leaders*. However, since Hotstuff reports a near-linear increment of latency with an increasing number of nodes, it only evaluated up to 128 nodes, where an attacker can DoS attack or partition one-third of all nodes rather than finding the leader. HoneyBadger [29] uses randomization to remove the partial synchrony assumption of PBFT. However, both its paper and our evaluation shows that HoneyBadger achieves high latency due to multiple rounds of its asynchronous byzantine agreements.

MinBFT [31] is an SGX-powered BFT protocol that has the same fault model as EGES. MinBFT reduces the number of rounds in PBFT and can tolerate more node failures (one-half instead of one-third), but MinBFT still has $O(n^2)$ message complexities, so its performance is not scalable to the number of nodes.

2.3 Consensus for Permissionless Blockchains

Existing permissionless blockchains can be divided into two categories based on how they confirm blocks. The first category confirms block with variants of the longestchain rule (i.e., Nakamoto consensus [1]), including Bit-Coin [1], Ethereum [27], BitCoin-NG [42], Snow-White [43], Ouroboros [44], Paros [45], Genesis [46], and GHOST [47]. Specifically, each node asynchronously selects the longest chain it received and confirms a block when there are kblocks succeeding it. However, waiting for k more blocks leads to a long confirm latency, and previous work [48] shows that this k must be large enough to ensure consistency. Moreover, the longest-chain rule cannot ensure consistency under partition attacks [49], [13], [11]. Intuitively, during a network partition, each partition will independently grow a chain; if these chains diverge for more than k blocks, nodes in different 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 [11], ByzCoin [18], Tendermint [17], and PeerCensus [50]. 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 [11], Tendermint [17], Tenderbake [51], and Tenderand [52]) 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 [53] and Tenderbake [51] 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 [8]) 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 [18] 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 [19] 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 SYSTEM MODEL

EGES is a consensus protocol for a permissioned blockchain running on M member nodes (*nodes* for short) connected with an asynchronous network. Each node is equipped with an attested SGX enclave running the EGES protocol (§5).

EGES adopts the hybrid fault model used in existing SGXpowered consensus protocols [54], [31], [55], 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. 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 two additional assumptions about the DoS attacker's capability as described below.

SGX's threat model. EGES has the same threat model for SGX as typical SGX-based systems [56], [57], [58], [23], [59]. 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.

Communication model. EGES maintains safety in an asynchronous network, where network packets can be dropped, delayed, or reordered arbitrarily. Nodes may be nonresponsive, either due to going offline or due to targeted DoS attacks (e.g., botnet DDoS attacks [15]) 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.

To achieve liveness, EGES has the "strong synchrony" assumption, same as Algorand [11]. Specifically, EGES assumes that messages between two nodes not under DoS attacks can be delivered within a known time bound, and the assumptions about the attacker's capability are described below.

Nodes are connected with a P2P overlay network, same as existing large-scale blockchain systems [1], [11]. Specifically, each node has a P2P module connecting to a random set of other nodes and relays messages using the gossip protocol [60].

A node's P2P module is outside SGX and can be controlled by the attackers: the attacker can partition some nodes from other nodes [12], [13]) or selectively pass consensus messages to nodes' SGX enclaves. However, such manipulations are already included in EGES's asynchronous network assumption. For safety, EGES leverages the sampling merit to estimate the delivery rate of a specific block proposal and derives overwhelming probability, regardless of how nodes are connected. For liveness, the adversary can control the P2P modules of a number of nodes with the restriction of the adversary's attack budget described below.

Capability of DoS attackers. EGES has three assumptions on the capability of a DoS attacker, same as Algorand [11] and existing move target defense (MTD) systems [25], [26]. First, the adversary has a targeted attack budget *B* (e.g., *B* = 300 or 10% of the total number of nodes): the adversary cannot constantly cause more than *B* targeted nodes in EGES to be nonresponsive. This budget is adaptive: the adversary can attack different nodes at different times, but the number of attacked nodes at a time cannot be constantly larger than B.

Second, the attacker can conduct ubiquitous DoS attacks (without targeting specific nodes) or partition a number of nodes from other nodes (e.g., by manipulating nodes' P2P modules [13], [12]). However, the P2P overlay network should have a large enough portion (e.g., 65%) of nodes connected. We provide a quantitative analysis of how EGES can preserve liveness under such attacks in §6.3.

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

EGES already assumes a strong enough attacker for practical distributed systems on the Internet. As pointed out by Algorand [11], a more powerful adversary than our model usually controls the internet service provider and can prevent all EGES nodes from communicating at all: no practical system can ensure liveness under such a strong adversary, and such attacks can be easily detected. We will provide a rigorous analysis of EGES's DoS resistance in §6.2.

4 EGES'S HIGH-LEVEL IDEA

EGES has three important features to achieve DoS resistance. First, EGES randomly selects a distinct group of committee for each block. The selection is done inside the SGX enclaves of a previous committee, and the selection result is encrypted on the confirmed blockchain. By doing so, a committee node can determine its committee membership without interactions with other nodes, making it *stay stealth* before trying to achieve consensus on its block.

Second, when a committee is achieving consensus for a given block, EGES uses fake committee nodes to conceal the real ones by sending dummy messages. Since whether a node is a real committee node is only known within the node's SGX enclave, and the dummy messages are of the same format as real ones, a DoS attacker cannot distinguish the real committee nodes from the fake ones. Therefore, the attacker must have an unrealistic large attack budget to attack all the real and fake committees; otherwise, he has to randomly guess who are the real ones.

Third, even if the attacker luckily guesses the real ones (and he may eventually succeed if tried persistently), EGES can ensure safety with overwhelming probability leveraging the delivery rate on all nodes of the unique proposal for each block. Specifically, even if a committee cannot confirm its own block, committees for subsequent blocks can help to consistently confirm this block by repetitive querying. This feature is in contrast to most existing consensus protocols (i.e., all except Algorand [11]), where the system must wait statically until a quorum of nodes become reachable.

EGES's consensus protocol has three parameters, n_A (default 300), τ (default 59%), and D (default 4), where n_A



Fig. 1: EGES's block status diagram.

is the number of acceptors, τ is the quorum ratio, and *D* is the finalize depth for an empty block. We will show how to select these parameters in §6.3 and how these parameters affect EGES's performance in §8.3.

For each block index n, EGES selects (§5.2) only one committee from all nodes. The committee contains two types of nodes: **one proposer** (P_n) and **a group of stealth acceptors** (\mathbb{A}_n) with the count of n_A . EGES ensures the following invariant (see §6.1 for proof).

Invariant 1. 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.

EGES uses different steps to confirm $proposal_n$ or $empty_n$. We make the steps of confirming $proposal_n$ as lightweight as possible for high efficiency in gracious runs.

Confirming proposal_n. In a gracious run, EGES confirms proposal_n in three steps: P_n broadcasts proposal_n with a propose request through the P2P network; acceptors send ACKs to P_n after receiving proposal_n; P_n broadcasts a finalize message confirming proposal_n once receiving ACKs from a quorum ($\tau \times n_A$) of acceptors.

Confirming empty_{*n*}. For consistency, when confirming $empty_n$, EGES must ensure that no node has confirmed proposal_{*n*}. Existing consensus protocols on static nodes use view-change protocols [61], [3], [17], [5] that count how many nodes have sent out ACKs and leverage the quorum intersection property [62] to determine whether a proposal may have been confirmed. However, in EGES, this method is not viable because EGES must ensure liveness even if most nodes in \mathbb{A}_n are DoS attacked after sending out their ACKs.

EGES's protocol for confirming $empty_n$ leverages the idea of repeated sampling to predicate that the probability of proposal_n having been confirmed is overwhelmingly low. 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; this is because confirming if proposal_n needs quorum ACKs from nodes in \mathbb{A}_n , and \mathbb{A}_n is uniformly selected from all nodes. Therefore, if one repetitively (iD) samples many nodes ($i\tau \times n_A$) from all nodes, and no node has received proposal_n, she 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 at the same time of seeking ACKs for their own proposals. Suppose P_n failed before sending out proposal_n, the next 4 proposers will all report that no node has received proposal_n, so empty_n can be confirmed as empty together with proposal_{n+1}~proposal_{n+4}.

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

First, Algorand uses verifiable random functions (VRF) to determine committees, so it can only control the expected count of proposers for each block without an exact number (1 \sim 70 in their experiment [11]). This design makes Algorand's consensus protocol not responsive [5], [8]: informally, a responsive protocol lets nodes wait for a number of messages rather than a large amount of time in each protocol step, which ensures a good performance when the network is in good condition. For each block, Algorand selects $1 \sim 70$ proposers, and each proposer broadcasts a block proposal with a distinct priority level. Then, Algorand selects one of these proposals by letting nodes vote for the received proposal with the highest priority. Since the total number of proposers is unknown, each node must wait for a conservatively long time (e.g., 10s) before voting to ensure it received most proposals. In contrast, EGES selects one proposer for each block, without the necessity for the selection progress, and EGES's protocol is responsive (§5).

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

Why does EGES use SGX? EGES chooses to use SGX for three main reasons. First, EGES uses SGX to regulate the behaviors of randomly selected committee nodes. As discussed in §2.3, using randomly selected committees in a permissioned blockchain may result in an unsolvable byzantine problem where more than one-third of committee nodes are malicious [9].

In EGES, each node have 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 (§7.1). A valid protocol message must carry a valid signature using the sender node's private key, proving that the message is generated in the sender node's enclave with code integrity. By doing so, a node cannot equivocate (i.e., sending conflicting messages to different other nodes) 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. 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 [11]). 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 discussed above, this makes EGES more efficient than Algorand's.

5 EGES CONSENSUS PROTOCOL

5.1 Protocol Preliminaries

Block structure. EGES adds one data field to the block structure of common blockchain systems [27], [1]: the encrypted committee identities for a future block (§5.2). EGES is oblivious to how transactions are stored or executed.

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 learntProposals, 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 $proposal_n$.

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 that, if a node's n^{th} is finalized as one choice, no other nodes will finalize the n^{th} block as the other choice. A node appends finalized proposal_n if it receives the finalize message for proposal_n; a node changes the empty_n from the undecided state to finalized if the node can predicate that no node has finalized proposal_n (§5.3).

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 [27], [10].

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 \mathbb{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 learntProposals is the set of known proposals for undecided blocks on this node and is saved in EGES enclave.

Membership and key managements. In EGES, 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 [3]: 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 only signs on message digests.

To ease understanding, in this section, we assume that EGES runs 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 will show how EGES supports dynamic membership and attestations in §7.1.

		Propose				
n		The index of the proposal	A	Igorithm 2: all nodes' action		
blk		the content of the block to be proposed	/	/* All message senders' memberships and signatures are		
MC_p		the MC value of the proposer		verified, ommited in all algorithms for brevity	*/	
\mathbb{U}_p		The \mathbb{U} list of the proposer	1	upon receiving msg = $\langle Propose, n, blk, MC_p \rangle_{\sigma_i}$		
pro	$posal_{um}$	$u_m = \max(\mathbb{U}_p)$. The proposer tries to finalize this	2	cache[n] = msg		
		proposal together with its proposed n^{th} block	3	if $MC < MC_p$: ask peers for missing blocks		
		ACK	4	if $hash(sk_i, n) > threshold$: Reply fake (cover) ACKs		
n		The index of corresponding proposal		message /* same format as true ACKs, with isReal =	+/	
sen	der	sender's public key (account) to identify it			-/	
not	ifications	proposals that the sender received and want to	5 1	upon receiving (Finalize, n, u_m , learnt) σ_i		
· D	1	notify the proposer	6	If $u_m \neq nu$:		
1SK6	eal	for identifying real ACK from cover messages	7	$[u_m]$.status \leftarrow intalized		
nor	ice	distable	8	$\begin{bmatrix} 0 & -0 \\ if \end{bmatrix} = 0 + Confirm chain up to index p (NC + p)$		
Tinaliza		9 If $ \cup == 0$: Confirm chain up to index n (MC \leftarrow n)				
		The index of the block finalized	10	learntProposals insert(learnt)		
n		The index of the block finalized	12	for u in \mathbb{I} in descending order \cdot		
loar	mt	Proposals learnt from acceptors potifications	13	$\int \mathbf{i} \mathbf{f} \operatorname{canFinalize}(u)$:		
lca	int	1 loposais learne from acceptors notifications	14	chain[u].status = finalized		
TABLE 2: ECES's concensus messages' fields. Blue fields		15	else: break // Empty blocks must be finalized in			
IAD		in the sheating weath and are left as within		descending order.		
are c	only used	In the checking mode and are left as nil in	16	upon timeout waiting for next block		
the r	normal m	ode.	17 chain.append (empty, status = undecided)			
A 1		t, the summary for the st th let t .	18 function can Finalize (u) :			
Algorithm 1: the proposer for the <i>n</i> th block		19	count = 0			
1 $n_A \leftarrow$ the number of acceptors		20	for $i = u + 1$; $i < chain.len; ++i$:			
2 blk \leftarrow the content of the n^{th} block to propose		21 if chain[i] != empty :				
3 function normalPropose():		22 count++				
4 bcast (Propose, blk, MC, nil, nil) $_{\sigma}$		23 if <i>count</i> $\geq D$: return true				
5 upon receiving (ACK, n, pk_i , nil) σ_{me}		24	else: return false			
6	ACKs.in	$nsert(pk_i)$				
7	if ACKs	s.count >= $\tau \times n_A$: bcast (Finalize, n, nil,				
	$\operatorname{nil} \sigma_n$	ne		the second se		
s function checkPropose():		Algorithm 3: an acceptor for the $n^{\iota n}$ block				
9	$u_m \leftarrow ma$	$\operatorname{ax}(\mathbb{U})$	1	upon receiving (Propose, blk, MC _n , \mathbb{U}_n , proposal) σ_i		
10	if cache[u ₁	m] != nil learntProposals[u_m] != nil :	2	r = rand()		
11 proposal $_{u_m} \leftarrow$ the proposal for index u_m		3 if $MC_p < MC$: notify proposer to catch up (omitted)				
12 else: $proposal_{u_m} \leftarrow nil$		4	$ \mathbf{if} \mathbb{U}_p == 0:$			
13 bcast (Propose, blk, MC, \mathbb{U} , proposal $_{u_m} angle_{\sigma_{me}}$		5	reply ({ACK, n, pk_{me} , nil, true, r} $_{pk_i}$) $_{\sigma_{me}}$			
14 $ \text{learnt} = []$			// Only the leader(i) can decrypt this message			
15 upon receiving (ACK, n, pk_i , notifications) σ_i		6	else:			
16 ACKs.insert(pk_i)		7	notifications = []			
17 learnt.insert(notifications)		8	$\int \mathbf{for} u \mathbf{n} \bigcup_p :$			
18	II ACKS	$f(timeline n memore) = \frac{1}{2} \frac{1}{$	9	If <i>cache[u]</i> != <i>nul</i> : notifications.append(cache[u])		
19		$(r inacize, n, proposal_{u_m}, icam)$	10	$[reply ({ACK, n, notifications, true, r}_{pk_i})_{\sigma_{me}}$		

Fig. 2: EGES's consensus protocol.

5.2 Selecting a Stealth Committee

For each block, EGES selects a committee, including **one proposer** and n_A **acceptors**, 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 (lookback) 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 lbcommittees' 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 (§5.1), 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 $(n_A + 1)$ 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, however, 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 knows an *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 §8.3).

5.3 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 $3\sim7$ 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 and its MC. The MC value helps nodes align confirmed parts of their chain: if a node's MC is smaller than the proposer's, the node asks for the missing confirmed blocks from its peers (Algorithm 2 Line 3). Upon receiving this propose request, an acceptor replies an ACK using UDP directly 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 which nodes are acceptors, but EGES's ensures that a non-acceptor cannot send valid ACKs (§3). Third, P_n broadcasts a finalize message; on receiving this message, a node finalizes proposal_n.

Checking mode. P_n is in the checking mode if it has undecided blocks (i.e., \mathbb{U} is non-empty), and its workflow is shown in Algorithm 1 Line 8~19. P_n checks the status of its undecided blocks and tries to finalize them (if possible) by adding additional fields to the propose message.

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 proposal_u, either from the propose messages from P_u from the notifications of other nodes, we call u a "known undecided" block; (2) otherwise we call u "unknown undecided". For each unknown 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 finalize message in order to let subsequent proposers finalize it; otherwise, P_n carries a message stating that most acceptors in \mathbb{A}_n never received proposal_n, and a node receiving this message finalize empty_n if the node received such messages form more than D consecutive proposers (Algorithm 2 Line 18).

For known undecided blocks, P_n helps to finalize only proposal_{u_m} where $u_m = max(\mathbb{U})$ and leaves other blocks for subsequent proposers. We will explain later that only finalizing u_m is essential for EGES's safety (Figure 5).

Figure 3 shows how a proposer P_{102} helps to finalize an undecided proposal₁₀₀. Suppose P_{100} failed just before broadcasting its finalize message. Therefore, the 100^{th} 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 4 shows 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 5 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 $proposal_{204}$ and confirm $empty_{200}$ (3) P_{205} and P_{206} are in partition B, so they timeout waiting for the finalize message for $proposal_{204}$ and mark the 204^{th} 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: nodes in partition A confirm $empty_{200}$, while nodes in partition B confirm proposal₂₀₀.

The inconsistency happens because, without this restriction, when a node finalizes $empty_n$, it only ensures proposers for blocks with index $\leq n + D$ has not finalized proposal_n;



Fig. 5: An example showing why a checking mode proposer can finalize only the block proposal with index = $max(\mathbb{U})$. D = 4 in this example.

adding this restriction helps to ensure that proposers for blocks with index > n + D cannot finalize proposal_n. Specifically, if this restriction presents in the previous example, nodes in partition B need to first finalize the empty₂₀₄ finalizing proposal₂₀₀. However, since proposal₂₀₄ has already been delivered to a large portion of nodes, this inconsistency cannot happen. We show the proof in §6.1.

5.4 Handling DoS Attacks Targeting Proposers

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

To mitigate this problem, we propose a new role of nodes called arbiter. An arbiter for a block index n does not generate new block proposals but only helps the proposer to finalize its proposal. For each block index n, EGES has many arbiters (larger than the attack budget B) doing the same tasks to tolerate DoS attacks on them.

Since we do not need to know the exact number of arbiters count and their identities, EGES simply let each node's EGES enclave independently determine whether it is an arbiter for the n^{th} block with a probability of p_a after receiving proposal_n. An arbiter for the n^{th} block broadcasts an arbit request following the same protocol as the proposer (Algorithm 1), and an acceptor responds to an arbit message with the same logic responding to a propose message.

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

6 SECURITY ANALYSIS

6.1 Safety with Overwhelming Probability

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

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

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

Proof. Proving this lemma is trivial if EGES works on a fixed member list, and we will show in §7.1 that EGES's protocol for dynamic memberships also ensures this lemma.

Lemma 2. (Invariant 1 in §4): at most one proposal can be generated for the n^{th} block.

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

Proof of the induction step. In EGES, each block has only two choices (Lemma 2), and a confirmed block must be first finalized (§5.1). Therefore, it is sufficient to prove the following proposition 1: the probability that one node finalizes $empty_n$ (event A), and another node finalizes $proposal_n$ (event B) is overwhelmingly small.

For event A, suppose node X finalizes $empty_n$. We use f_{m_x} to denote the maximum finalized block index on node

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

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

$$Prob(B1) = Prob(\text{Re} > \tau \times n_A)$$

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

Therefore, the probability of event A is

$$Prob(A) = (Prob(NR > \tau \times n_A))^D$$

The calculation shows that the probability of event A and event B1 happening together $Prob(A) \times Prob(B1)$ is overwhelmingly low for any delivery rate p by setting τ and n_A (§8.3). For instance, our evaluation chose τ as 59%, D as 4, n_A as 300, M as 10K, and the probability of EGES enforcing safety is $1-10^{-9}$. In real deployments, M may change due to membership changes; however, when M is much bigger (e.g., 20X) than n_A , this probability is not sensitive to M because hypergeometric distribution is approximate to binomial.

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

Putting the two steps together, we proved proposition 1 and thus proved the induction step that the n^{th} block must be confirmed consistently among nodes with overwhelmingly high probability. Therefore, EGES ensures safety with overwhelmingly high probability.

6.2 Liveness under Targeted DoS Attacks

EGES can defend against DoS attacks and partition attacks targeting committee nodes under the threat model in §3. Since from a single node's point of view, when it cannot reach a remote node, it cannot distinguish whether a remote node is under DoS attack or is partitioned, EGES handles these two attacks altogether.

EGES has two types of committee nodes for each block, a proposer and n_A acceptors, randomly selected from all nodes in the system. Because of EGES's stealth committee abstraction (§5.2), the identity of each committee node is unknown to attackers outside SGX enclaves before the node sending out its first protocol message.

We first discuss acceptors. An acceptor sends ACK messages to both the proposer and arbiters, but its identity becomes explicit after sending out its first ACK message, so EGES uses fake acceptors to conceal the real ones. If observed from outside enclaves, the fake acceptors behave identically as real ones so that an attacker cannot differentiate the real acceptors and attack them. EGES achieves this with three design points.

First, fake acceptors are randomly selected from all nodes for each block so that an attacker cannot determine the fake acceptors by monitoring network packets (§5.2). Second, real and fake acceptors' EGES enclaves respond to protocol messages (propose and arbit) in the same way if observed outside their enclaves (§5.3), so that an attacker cannot distinguish real or fake acceptors by watching their behaviors. Third, all messages from real or fake acceptors are of the same format, encrypted with the receiver's public key (5.3) and can only be decrypted in the receiver's enclave, so that an attacker cannot differentiate real acceptors from fake ones by watching the packet content.

Therefore, if an attacker targeting acceptors in EGES, it can only randomly select *B* nodes from both real acceptors and fake acceptors. For instance, if EGES has $n_A = 300$, $\tau = 59\%$ and has (expected) 600 fake acceptors; if the attacker's attack budget is 300, the probability that the attacker can luckily attack more than $(1 - \tau) \times n_A$ real acceptors for one block is 0.3%. Moreover, even if the attacker is so lucky that it successfully guessed more than $(1 - \tau) \times n_A$ acceptors for some block (say n^{th}), EGES can still consistently determine whether to confirm proposal_n or empty_n using the committees for subsequent blocks (§5.3); in other words, the attacker must constantly be so lucky to make EGES lose liveness.

Then we discuss proposers. The identity of a proposer (say P_n) becomes explicit after broadcasting its proposal_n and may be targeted attacked when it is waiting for quorum ACKs. However, since EGES has many (larger than B) arbiters that can help to finalize the proposal_n, attacking P_n will not affect EGES's liveness. EGES lets each node independently determine whether it is an acceptor for each block (§5.4) so that an attacker cannot keep attacking them. Note that one single proposer or arbiter is enough for EGES to stay alive because any of them can finalize the current n^{th} block.

Note that EGES's targeted attack model (§3) handles only DoS or partition attacks targeting specific EGES nodes. A more powerful attacker may also target EGES's major communication links. From the protocol aspect, EGES avoids such vulnerabilities in the Network layer (in the OSI model [63]) by using distinct committees for different blocks: EGES's protocol traffic is spread among the whole P2P network rather





Fig. 6: Parameter selection for τ and *D* on 10K nodes for different n_A values.

Fig. 7: Connected component size required to ensure liveness with different *D* values.

than centralized among a few dedicated nodes. However, when EGES is deployed, EGES's communication messages may be aggregated in the Link or Physical layer. For instance, if a large number of EGES nodes are hosted in the same data center (DC), the links connecting this DC and the Internet may be susceptible to attacks. However, such attacks are not adaptive, and as long as a great majority of nodes are connected, EGES can achieve practical liveness. §8.3 shows the relation between EGES's liveness and the maximum connected component size in the P2P network. Nevertheless, EGES cannot ensure liveness under arbitrary partitions, and previous work shows that it is impossible to ensure both consistency and liveness under partitions [24], [64].

6.3 Parameter Selection

Figure 6 shows the relation between τ and D in order to ensure 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 (§5.3).

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 7. 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 (§8.2).

Figure 8 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 (§6.1) becomes a binomial distribution with the probability of $p \times M/n_A$, and other distribution changes similarly. As shown in Figure 8, EGES would need a larger quorum size $\tau \times n_A$, and worse liveness on network partition.

7 IMPLEMENTATION

We selected the Golang implementation of Ethereum (i.e., geth) as our codebase because geth is heavily tested on



Fig. 8: Parameter selection and liveness requirements if EGES lets each node to independently decide its committee membership.

the Internet. We leveraged the P2P libraries from geth and rewrote the functions for generating, verifying, and handling new blocks. Since SGX only provides SDKs in C/C++, we used CGo to invoke ECalls. We modified 2073 lines of Golang code and implemented the consensus protocol for 1943 lines of C code. For asymmetric key based encryption, we used ECC-256 from the API provided by the SGX SDK.

Each EGES node has three modules: a consensus module running the EGES consensus protocol and storing nodes' member list (§5), a P2P module connecting to a random set of peers and relaying messages using the Gossip [60] protocol, and a blockchain core module storing confirmed parts of chain and transactions submitted by clients.

As stated in §3, only the consensus module runs in the node's SGX enclave. We put the P2P module outside SGX because this module is only responsible for relaying messages and putting it inside SGX cannot effectively bring more benefits: a malicious node can still selectively drop packets at the network layer. We put the blockchain core module outside enclave because immutable properties of the hash-connected blockchain make this module's actions easily verifiable, and putting this module in enclave will consume much memory.

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 lefts the block's field of "hash value for the previous block" empty, and EGES's enclave computes this field when confirming this block. In other words, EGES actually achieves consensus on a totally ordered sequence of transactions, same as Hyperledger Fabric [10], and encapsulate these baches into a hash-chain of blocks while confirming them.

7.1 Membership and attestation

To support dynamic memberships, EGES leverages the idea from SCIFER [35] to record the joining of new nodes as transactions on the blockchain. This mechanism ensures Lemma 1 because all updates to the member list are only determined by confirmed blocks.

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

A node *i* joins EGES with three steps. First, *i* launches its EGES enclave, which generates its account (pk_i, sk_i) and creates the hardware monotonic counter *c* for defending forking attacks (§7.3). The node's account is securely saved to permanent storage using SGX's seal mechanism [22] for recoveries from machine failures (e.g., power off). Then, *i* sends a *join* request to *j*. Second, *j* does a standard SGX



(c) On receiving a propose request

Fig. 9: EGES's enclave interactions (ECalls and OCalls). The enclave module is shaded in orange.

remote attestation [22], which succeeds with a signed quote Q_i from Intel's attestation service, and *i*'s enclave transfers its public key pk_i and counter value *c* to *j*'s enclave through the secure communication channel between two enclaves created during attestation. Third, node *j*'s enclave creates a signed registration transaction including pk_i , *c*, $addr_i$, Q_i and *i*'s ip address. Node *i* joins EGES when the transaction is included in a confirmed block.

7.2 Enclave Interactions

Figure 9 shows the implementation of EGES enclave. An EGES enclave holds three data structures shared among ECalls: cache, is_proposer, and is_acceptor, each as a hash map. As explained in §5.1, the cache saves received block proposals. In our implementation, the cache only keeps the hash values of block headers instead of whole blocks to save enclave memory. is_proposer and is_acceptor are hash maps with block indices as keys and boolean as values, saving whether this node will be in the committee for a future block.

In Figure 9a, when a node's blockchain core module confirms the $(n - lb)^{th}$ block, it asynchronously invokes an ECall letting the enclave check whether it will be the proposer/acceptor for the n^{th} block (§5.2). In Figure 9b, when a node's core module appends the $(n - 1)^{th}$ block, it 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 9c, 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 (§5.4).

7.3 Handling Attacks on Enclaves

Forking attacks. In the P2P scenario, one challenge is enclave forking attacks [65]. EGES must permit a node to reuse its sealed account (§7.1) 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 [65], [66] 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 [22], which is monotonic among all enclaves on the same machine. When a node launches its EGES enclave, the enclave increments and read this counter value c and enclose 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.

Premature timeout. EGES uses the timeout mechanism to suspect node failures and to achieve liveness. A premature timeout does not affect EGES's safety because EGES's safety argument does not rely on timing constraints (§6.1).

However, a premature timeout may affect EGES's performance, and we implemented a simple mechanism to mitigate this problem by leveraging the trusted timer API *sgx_get_trusted_time* provided by the SGX platform. This API provides a verifiable timer with the granularity of seconds. To avoid the heavy time cost (tens of milliseconds) caused by frequently checking this timer, EGES maintains an untrusted timer outside. When the untrusted timeout is triggered, the untrusted code invokes the timeout logic in EGES's consensus module in the enclave, and the consensus module checks the trusted timer before invoking EGES's consensus logic.

8 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 typical georeplicated setting, while running EGES on both our cluster and AWS, we emulated the world scale Internet by using the Linux traffic control (TC) command 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

Config	Cluster	AWS Cloud
# Nodes	300	up to 10K
Acceptor group size (n_A)	100	300
D	4	4
au	65%	58%
LB	5000	10000
timeout	2s	3s
SGX mode	hardware mode	simulation mode

TABLE 3: EGES's evaluation parameters.

hardware mode because EGES's performance is bound to network latency in WAN (§8.1). The scalability (Figure 10) and robustness (Figure 11) 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 [28], SBFT [4], HoneyBadger [29], and HotStuff [5]), two SGXpowered consensus protocols for permissioned blockchains (Intel-PoET [30] and MinBFT [31]), the default consensus protocol in our codebase (Ethereum-PoW [27]), and two permissionless blockchains' protocol that runs on dynamic committees (Algorand [11] and Tendermint [17]). A detailed description of these protocols is in §2.2.

Since Algorand's open-source code is still under development, and we were unable to deploy its latest release [67] to the same scale as EGES and Algorand's own paper (i.e., 10k nodes), we took Algorand's performance from Figure 5 in its paper [11]. 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 clientside 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 [30]), 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 [68], [53]. 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:

§8.1: Is EGES efficient and scalable?

§8.2: What is EGES's performance under targeted DoS attacks?

§8.3: How sensitive is EGES to its parameters?



Fig. 10: Scalability to the number of nodes on Internet.

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

§8.5: What are the limitations and potential future works of EGES?

8.1 Efficiency and Scalability

Table 1 shows the performance comparison of EGES and eight baseline protocols. As Alogrand's paper [11] 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. §8.4 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 allto-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 causes a dramatic performance drop to the system, which is evaluated in §8.4.

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 [69].

Breakdown and micro-events. To understand EGES's latency, we recorded the time taken for the two steps of EGES's protocol (§5.3): 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 takes around 3us [22], 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.

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.

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 10 shows the result. The latency is divided into two parts. The figure shows that the seeking for quorum ACKs phase of EGES (§5.3) 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 [70], 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. While running on AWS, EGES's latency was slightly faster than on our cluster, because AWS CPUs are faster.

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). 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 (§5.3).

8.2 Performance on DoS Attacks

To evaluate EGES's robustness under DoS attacks, we ran EGES with 1000 nodes on AWS with $n_A = 100$, and conducted targeted DoS attacks that are compliant with our attack model (§3): 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 and 50 correspondingly. 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 §8.3, 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 execution log and found that the slow recovery was because the proposer for the next block happened to be



Fig. 11: EGES's throughput on DoS and network partition attacks. There were 1000 nodes on AWS at 0s.

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 (§5.4). 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 11b shows the throughput measured in the large partition. Overall, the large partition maintained liveness during the partition. The small partition 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. §8.3 shows a quantitative analysis of how EGES can preserve liveness under network partitions.

8.3 Sensitivity

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 13 shows 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.

Figure 12 shows 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.

8.4 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 14a 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



Fig. 12: EGES's sensitivity on block sizes (cluster setting).



Fig. 13: EGES's sensitivity on acceptor size and expected fake acceptor numbers (cluster setting).

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 14a also shows SBFT's performance. In the nongeo-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 14b 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's 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 straggler nodes (usually two (§2.2)), we



Fig. 14: Comparing EGES, SBFT, and BFT-SMaRt.

scenarios, including DoS attacks (Figure 11a). For comparison, we evaluated the performance of BFT-

reverted to the slow path (PBFT).

comparison, we evaluated the performance of BFT-SMaRt and SBFT on node failures (i.e., DoS attacks targeting consensus nodes). Figure 14c 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 14d), 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.

observed that 87% of the consensus rounds in SBFT have

More importantly, EGES can safely switch its acceptor groups across blocks, and it tolerated various failure

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.

8.5 Discussions

EGES has two limitations. First, EGES requires a joining 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 public blockchains [21] and permissioned blockchains [30], [20], [31] also use SGX. Second, EGES targets large-scale, geo-distributed permissioned blockchain systems (e.g., a global payment system [8]), while for small-scale deployments (e.g., supply chain among a few small companies) in a single data center, existing consensus protocols (e.g., BFT-SMaRt) are more suitable.

EGES can incite the deployment of more Internet-scale applications that highly desire the DoS resistance feature of EGES, including e-voting, decentralized auction, and payment systems. Moreover, the attested SGX enclaves on EGES nodes bring the potential to port existing centralized SGX-powered applications on to EGES, including SGX-protected distributed database [56], [57], [71] and SGX-powered anonymous networks [72]. For instance, one can deploy an EGES-ToR application by letting the EGES enclave on each node do a local attestation for an SGX-ToR enclave [72] and letting EGES nodes maintain the ToR directory service on the blockchain. By doing so, EGES-ToR removes SGX-ToR's centralized directory service that is vulnerable to DoS attacks or service censoring.

9 CONCLUSION

We have presented EGES, the first efficient permissioned blockchain consensus protocol that can tolerate targeted DoS and partition attack. EGES achieves comparable performance to existing permissioned blockchain's consensus protocols while achieving much stronger robustness. EGES is carefully implemented with two promising distributed applications, greatly improving the reliability and security of their legacy, centralized versions. EGES's source code is available on github.com/hku-systems/eges.

REFERENCES

- S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," https://bitcoin.org/bitcoin.pdf, Dec 2008, accessed: 2015-07-01. [Online]. Available: https://bitcoin.org/bitcoin.pdf
- [2] J. Sousa, A. Bessani, and M. Vukolić, "A byzantine fault-tolerant ordering service for the hyperledger fabric blockchain platform," IEEE/IFIP International Conference on Dependable System and Network (DSN 2018), 2017, accessed:2017-09-25. [Online]. Available: https://arxiv.org/pdf/1709.06921.pdf
- [3] M. Castro and B. Liskov, "Practical byzantine fault tolerance," in Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI '99), Oct. 1999.
- [4] G. G. Gueta, I. Abraham, S. Grossman, D. Malkhi, B. Pinkas, M. K. Reiter, D.-A. Seredinschi, O. Tamir, and A. Tomescu, "Sbft: a scalable decentralized trust infrastructure for blockchains," IEEE/IFIP International Conference on Dependable System and Network (DSN 2019), 2019. [Online]. Available: https: //arxiv.org/pdf/1804.01626.pdf
- [5] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hotstuff: Bft consensus with linearity and responsiveness," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM, 2019, pp. 347–356.
- [6] "Medical Chain," http://www.medicalchain.org/, 2017.
- [7] "Blockchain for Supply Chain," https://www.ibm.com/ blockchain/industries/supply-chain, 2017.
- [8] M. Baudet, A. Ching, A. Chursin, G. Danezis, F. Garillot, Z. Li, D. Malkhi, O. Naor, D. Perelman, and A. Sonnino, "State machine replication in the libra blockchain," 2019.
- [9] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," vol. 4, no. 3. ACM, 1982, pp. 382– 401. [Online]. Available: http://people.cs.uchicago.edu/~shanlu/ teaching/33100_wi15/papers/byz.pdf
- [10] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich et al., "Hyperledger fabric: a distributed operating system for permissioned blockchains," in *Proceedings of the Thirteenth EuroSys Conference (EuroSys 2018)*. ACM, 2018, p. 30.
- [11] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," Cryptology ePrint Archive, Report 2017/454, 2017, accessed: 2017-06-29. [Online]. Available: http://eprint.iacr.org/2017/454.pdf
 [12] C. Natoli and V. Gramoli, "The balance attack against
- [12] C. Natoli and V. Gramoli, "The balance attack against proof-of-work blockchains: The R3 testbed as an example," http://arxiv.org/abs/1612.09426, 2016, accessed: 2017-02-15. [Online]. Available: https://arxiv.org/pdf/1612.09426.pdf
- [13] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on bitcoin's peer-to-peer network," in 24th USENIX Security Symposium (USENIX Security 15), 2015, pp. 129– 144. [Online]. Available: https://www.usenix.org/system/files/ conference/usenixsecurity15/sec15-paper-heilman.pdf
- [14] "Banking in the United States," BankingintheUnitedStates.
- [15] O. Kupreev, E. Badovskaya, and A. Gutnikov, "Ddos attacks in q2 2019," 2019.
- [16] A. Rayome, "Major ddos attack lasts 297 hours, as botnets bombard businesses," 2018.
- [17] E. Buchman, "Tendermint: Byzantine fault tolerance in the age of blockchains," http://atrium.lib.uoguelph.ca/xmlui/bitstream/ handle/10214/9769/Buchman_Ethan_201606_MAsc.pdf, Jun 2016, accessed: 2017-02-06. [Online]. Available: http://atrium.lib.uoguelph.ca/xmlui/bitstream/handle/ 10214/9769/Buchman_Ethan_201606_MAsc.pdf
- [18] E. K. Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford, "Enhancing bitcoin security and performance with strong consistency via collective signing," in 25th USENIX Security Symposium (USENIX Security 16). Austin, TX: USENIX Association, Aug. 2016. [Online]. Available: http://arxiv.org/pdf/1602.06997.pdf
- http://arxiv.org/pdf/1602.06997.pdf
 [19] P. Jovanovic, "Byzcoin: Securely scaling blockchains," http://hackingdistributed.com/2016/08/04/byzcoin/, 2016, accessed: 2019-08-01. [Online]. Available: http://hackingdistributed.com/2016/08/04/byzcoin/
- building [20] "Ccf: framework confidential А for verifiable replicated services," Microsoft, Tech. Rep. MSR-TR-2019-16, 2019. April [Online]. Available: https://www.microsoft.com/en-us/research/publication/ ccf-a-framework-for-building-confidential-verifiable-replicated-services/

- [21] F. Zhang, I. Eyal, R. Escriva, A. Juels, and R. van Renesse, "Rem: Resource-efficient mining for blockchains," http://eprint.iacr.org/2017/179, 2017. [Online]. Available: http: //eprint.iacr.org/2017/179.pdf
- [22] V. Costan and S. Devadas, "Intel sgx explained." IACR Cryptology ePrint Archive, vol. 2016, p. 86, 2016.
- [23] "hyperledger-labs/minbft," https://github.com/ hyperledger-labs/minbft.
- [24] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," vol. 32, no. 2. ACM, 1985, pp. 374–382. [Online]. Available: http://macs.citadel. edu/rudolphg/csci604/ImpossibilityofConsensus.pdf
- [25] J.-H. Cho, D. P. Sharma, H. Alavizadeh, S. Yoon, N. Ben-Asher, T. J. Moore, D. S. Kim, H. Lim, and F. F. Nelson, "Toward proactive, adaptive defense: A survey on moving target defense," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 1, pp. 709–745, 2020.
- [26] S. Venkatesan, M. Albanese, K. Amin, S. Jajodia, and M. Wright, "A moving target defense approach to mitigate ddos attacks against proxy-based architectures," in 2016 IEEE conference on communications and network security (CNS). IEEE, 2016, pp. 198–206.
- [27] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," https://github.com/ ethereum/wiki/wiki/White-Paper, 2014, accessed: 2016-08-22. [Online]. Available: https://github.com/ethereum/wiki/wiki/ White-Paper
- [28] J. Sousa, A. Bessani, and M. Vukolić, "A byzantine fault-tolerant ordering service for the hyperledger fabric blockchain platform," arXiv:1709.06921, 2017, accessed:2017-09-25. [Online]. Available: https://arxiv.org/pdf/1709.06921.pdf
- [29] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, "The honey badger of bft protocols," https://eprint.iacr.org/2016/199.pdf, 2016, accessed: 2017-01-10. [Online]. Available: https://eprint.iacr. org/2016/199.pdf
- [30] https://www.hyperledger.org/projects/sawtooth.
- [31] G. S. Veronese, M. Correia, A. N. Bessani, L. C. Lung, and P. Verissimo, "Efficient byzantine faulttolerance," vol. 62, no. 1. IEEE, 2013, pp. 16–30. [Online]. Available: https://www.researchgate.net/profile/ Miguel_Correia3/publication/260585535_Efficient_Byzantine_ Fault-Tolerance/links/5419615d0cf25ebee9885215.pdf
- [32] R. Riemann and S. Grumbach, "Distributed protocols at the rescue for trustworthy online voting," arXiv:1705.04480, 2017, accessed: 2017-06-29. [Online]. Available: https://arxiv.org/pdf/1705.04480. pdf
- [33] J. Aumasson and L. Merino, "Sgx secure enclaves in practicesecurity and crypto review," *Black Hat*, 2016.
- [34] M. Hamburg, P. Kocher, and M. E. Marson, "Analysis of intel's ivy bridge digital random number generator," Online: http://www. cryptography. com/public/pdf/Intel_TRN G_Report_20120312. pdf, 2012.
- [35] M. Ahmed and K. Kostiainen, "Identity aging: Efficient blockchain consensus," arXiv preprint arXiv:1804.07391, 2018.
- [36] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, and D. Song, "Ekiden: A platform for confidentialitypreserving, trustworthy, and performant smart contract execution," arXiv preprint arXiv:1804.05141, 2018.
- [37] R. Yuan, Y.-B. Xia, H.-B. Chen, B.-Y. Zang, and J. Xie, "Shadoweth: Private smart contract on public blockchain," *Journal of Computer Science and Technology*, vol. 33, no. 3, pp. 542–556, 2018.
- [38] J. Lind, O. Naor, I. Eyal, F. Kelbert, E. G. Sirer, and P. Pietzuch, "Teechain: A secure payment network with asynchronous blockchain access," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, ser. SOSP '19. New York, NY, USA: ACM, 2019, pp. 63–79. [Online]. Available: http://doi.acm.org/10.1145/3341301.3359627
- [39] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi, "Town crier: An authenticated data feed for smart contracts," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer* and Communications Security. ACM, 2016, pp. 270–282. [Online]. Available: https://eprint.iacr.org/2016/168.pdf
- [40] I. Bentov, Y. Ji, F. Zhang, Y. Li, X. Zhao, L. Breidenbach, P. Daian, and A. Juels, "Tesseract: Real-time cryptocurrency exchange using trusted hardware," Cryptology ePrint Archive, Report 2017/1153, 2017, accessed:2017-12-04. [Online]. Available: https://eprint.iacr.org/2017/1153.pdf
- [41] M. Tran, L. Luu, M. S. Kang, I. Bentov, and P. Saxena, "Obscuro: A es/ bitcoin mixer using trusted execution environments," Cryptology

ePrint Archive, Report 2017/974, 2017, accessed:2017-10-06. [Online]. Available: http://eprint.iacr.org/2017/974.pdf

- [42] I. Eyal, A. E. Gencer, E. G. Sirer, and R. van Renesse, "Bitcoin-ng: A scalable blockchain protocol," in 13th USENIX Security Symposium on Networked Systems Design and Implementation (NSDI'16). USENIX Association, Mar 2016. [Online]. Available: http://www.usenix. org/system/files/conference/nsdi16/nsdi16-paper-eyal.pdf
- [43] I. Bentov, R. Pass, and E. Shi, "Snow white: Provably secure proofs of stake," https://eprint.iacr.org/2016/919.pdf, 2016, accessed: 2016-11-08. [Online]. Available: https://eprint.iacr.org/2016/919. pdf
- [44] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-ofstake blockchain protocol," https://pdfs.semanticscholar.org/ 1c14/549f7ba7d6a000d79a7d12255eb11113e6fa.pdf, 2016, accessed: 2017-02-20. [Online]. Available: https://pdfs.semanticscholar.org/ 1c14/549f7ba7d6a000d79a7d12255eb11113e6fa.pdf
- [45] B. David, P. Gazi, A. Kiayias, and A. Russell, "Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake protocol," Cryptology ePrint Archive, Report 2017/573, 2017, accessed: 2017-06-29. [Online]. Available: http://eprint.iacr.org/2017/573.pdf
- [46] C. Badertscher, P. Gazi, A. Kiayias, A. Russell, and V. Zikas, "Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security.* ACM, 2018, pp. 913–930.
- [47] Y. Sompolinsky and A. Zohar, "Accelerating bitcoin's transaction processing. fast money grows on trees, not chains," p. 881, 2013. [Online]. Available: http://eprint.iacr.org/2013/881.pdf
- [48] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," https://eprint.iacr.org/2016/555.pdf, 2016, accessed: 2016-08-10. [Online]. Available: https://eprint.iacr.org/2016/555. pdf
- [49] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking bitcoin: Routing attacks on cryptocurrencies," in 2017 IEEE Symposium on Security and Privacy (SP). IEEE, 2017, pp. 375–392.
- [50] C. Decker, J. Seidel, and R. Wattenhofer, "Bitcoin meets strong consistency," in *Proceedings of the 17th International Conference on Distributed Computing and Networking*. ACM, 2016, p. 13.
- [51] L. Aştefanoaei, P. Chambart, A. Del Pozzo, E. Tate, S. Tucci, and E. Zălinescu, "Tenderbake–classical bft style consensus for public blockchains," arXiv preprint arXiv:2001.11965, 2020.
- [52] S. Pyo Kim, "Tenderand: Randomized leader election in tendermint," 2020. [Online]. Available: https://medium.com/codechain/ tenderand-randomized-leader-election-in-tendermint-a3663d863479
- [53] Y. Amoussou-Guenou, A. Del Pozzo, M. Potop-Butucaru, and S. Tucci-Piergiovanni, "Correctness of tendermint-core blockchains," in 22nd International Conference on Principles of Distributed Systems (OPODIS 2018). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [54] B.-G. Chun, P. Maniatis, S. Shenker, and J. Kubiatowicz, "Attested append-only memory: Making adversaries stick to their word," in ACM SIGOPS Operating Systems Review, vol. 41, no. 6. ACM, 2007, pp. 189–204. [Online]. Available: http://news.cs.nyu.edu/~jinyang/fa08/papers/a2m.pdf
- [55] Q. Zhang, Z. Qi, X. Liu, T. Sun, and K. Lei, "Research and application of bft algorithms based on the hybrid fault model," in 2018 1st IEEE International Conference on Hot Information-Centric Networking (HotICN). IEEE, 2018, pp. 114–120.
- [56] F. Shaon, M. Kantarcioglu, Z. Lin, and L. Khan, "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)*, 2017.
- [57] C. Priebe, K. Vaswani, and M. Costa, "Enclavedb: A secure database using sgx," in *Proceedings of the 2018 IEEE Symposium on Security* and Privacy. IEEE, 2018, p. 0.
- [58] J. Lind, C. Priebe, D. Muthukumaran, D. O'Keeffe, P.-L. Aublin, F. Kelbert, T. Reiher, D. Goltzsche, D. Eyers, R. Kapitza et al., "Glamdring: Automatic application partitioning for intel sgx," in 2017 USENIX Annual Technical Conference (USENIX ATC 17), Santa Clara, CA, 2017.
- [59] J. Jiang, X. Chen, T.-O. Li, C. Wang, T. Shen, S. Zhao, H. Cui, C.-L. Wang, and F. Zhang, "Uranus: Simple, efficient sgx programming and its applications," in *Proceedings of the 15th ACM on Asia Conference on Computer and Communications Security 2020 (ASIACCS*)

'20, accepted). https://hemingcui.github.io/accepted/asiaccs20-uranus.pdf, 2020.

- [60] A.-M. Kermarrec and M. Van Steen, "Gossiping in distributed systems," ACM SIGOPS Operating Systems Review, vol. 41, no. 5, pp. 2–7, 2007.
- [61] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in *Proceedings of the USENIX Annual Technical Conference (USENIX '14)*, Jun. 2014.
- [62] H. Howard, D. Malkhi, and A. Spiegelman, "Flexible paxos: Quorum intersection revisited," arXiv preprint arXiv:1608.06696, 2016.
- [63] "OSI model Wikipedia," https://en.wikipedia.org/wiki/OSI_ model.
- [64] S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," vol. 33, no. 2. ACM, 2002, pp. 51–59. [Online]. Available: http://www.comp.nus. edu.sg/~gilbert/pubs/BrewersConjecture-SigAct.pdf
- [65] S. Matetic, M. Ahmed, K. Kostiainen, A. Dhar, D. Sommer, A. Gervais, A. Juels, and S. Capkun, "Rote: Rollback protection for trusted execution." *IACR Cryptology ePrint Archive*, vol. 2017, p. 48, 2017.
- [66] J. Gu, Z. Hua, Y. Xia, H. Chen, B. Zang, H. Guan, and J. Li, "Secure live migration of sgx enclaves on untrusted cloud," in 2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). IEEE, 2017, pp. 225–236.
- [67] "Algorand/go-algorand," https://github.com/algorand/ go-algorand/releases/tag/v2.0.14-beta.
- [68] "Cryptocurrency statistics," 2019. [Online]. Available: https: //bitinfocharts.com/
- [69] Deloitte, "Blockchain Performance Report," https://en.bitcoin.it/ wiki/Satoshi_Client_Node_Discovery#DNS_Addresses, 2018.
- [70] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," *SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 4, pp. 149–160, 2001.
- [71] T. Kim, J. Park, J. Woo, S. Jeon, and J. Huh, "Shieldstore: Shielded inmemory key-value storage with sgx," in *Proceedings of the Fourteenth EuroSys Conference* 2019. ACM, 2019, p. 14.
- [72] S. M. Kim, J. Han, J. Ha, T. Kim, and D. Han, "Enhancing security and privacy of tor's ecosystem by using trusted execution environments." in NSDI, 2017, pp. 145–161.



Xusheng Chen received his Bachelor degree in HKU. He is currently a PhD student in Computer Science of HKU (2017-now). He is under the supervision of Dr. Heming Cui. His research interests include distributed consensus protocols, distributed systems and system security.



Shixiong Zhao received his Bachelor degree in HKU and his master degree in HKUST. He is currently a PhD student in Computer Science of HKU (2017-now). He is under the supervision of Dr. Heming Cui. His research interests include distributed systems for high performance computing, distributed systems and system security.



Ji Qi received his B.S (2015) degree from Beijing Institude of Technology, Beijing, China, and his M.S (2018) degree from Tsinghua University, Beijin, China. He is currently pursuing the PhD in computer science at the University of Hong Kong under the supervision of Dr. Heming Cui. His interests include domain-specific modeling, distributed system and cloud computing.



Jianyu Jiang is currently a third year PhD student in Computer Science Department at The University of Hong Kong. He is working on topics in large scale computation platform under the supervision of Dr. Heming Cui. Jianyu received his Bachelor's Degree in Computer Science Department at Xi'an Jiaotong University, under the supervision of Professor Qi Yong.



Haoze Song received the BS degree from Department of Computer Science, University of Science and Technology of China, in 2020. He is currently working towards the MPhil degree in Computer Science at HKU. His research interests mainly focus on distributed system and distributed computing.



Cheng Wang received his PhD from the University of Hong Kong and B.Eng degree from Shanghai Tongji University. His research interests lie in distributed systems, with a particular focus on fault tolerance. He is currently working in Huawei Ltd.



Tsz On Li received his Bachelor degree in HKU. He is currently an MPhil student in Computer Science of HKU (2019-now). He is under the supervision of Dr. Heming Cui. His research interests include differential privacy and big data systems.



T-H. Hubert Chan T-H. Hubert Chan is an Associate Professor at the Department of Computer Science at the University of Hong Kong. He completed his PhD in Computer Science at Carnegie Mellon University in 2007. His main research interests are approximation algorithms, discrete metric space, privacy and security inspired problems.



Fengwei Zhang is an Associate Professor at Department of Computer Science and Engineering at Southern University of Science and Technology (SUSTech). His primary research interests are in the areas of systems security, with a focus on trustworthy execution, hardware-assisted security, debugging transparency, transportation security, and plausible deniability encryption. Before joining SUSTech, he spent four wonderful years as an Assistant Professor at Department t Wayne State University

of Computer Science at Wayne State University.



Xiapu Luo received his B.S. in Communication Engineering and M.S. in Communications and Information Systems from Wuhan University. He obtained his Ph.D. degree in Computer Science from the Hong Kong Polytechnic University, under the supervision of Prof. Rocky K.C. Chang. After that, he spent two years at the Georgia Institute of Technology as a post-doctoral research fellow advised by Prof. Wenke Lee. His current research interests include Network and System Security,

Information Privacy, Internet Measurement, Cloud Computing, and Mobile Networks.



Sen Wang received the B.S. degree from the University of Science and Technology of China (USTC) in 2005, the M.S. degree from the Chinese Academy of Sciences (CAS) in 2008, and the Ph.D. degree from Tsinghua University in 2014, all in computer science. From 2014 to 2019, he was a lecturer and then an associate professor at Chongqing University, China. Currently, he is a senior researcher at Huawei, Hongkong. His research interests include information-centric

networking, Federated Learning and AI for System.





Gong Zhang is a chief architect researcher scientist, director of the Huawei Future Network Theory Lab. His major research directions are network architecture and large-scale distributed systems. He has abundant experience on system architect in networks, distributed system and communication system for more than 20 years. He has more than 90 global patents.

Heming Cui is an associate professor in computer science of HKU. His research interests include operating systems, programming languages, distributed systems, and cloud computing, with a particular focus on building software infrastructures and tools to improve reliability and security of real-world software. Homepage: https://i.cs.hku.hk/ heming/. He is a member of IEEE.