

Robust P2P Primitives Using SGX Enclaves

Yaoqi Jia¹ Shruti Tople² Tarik Moataz³ Deli Gong¹ Prateek Saxena⁴ Zhenkai Liang⁴

¹ACM Member {jiayaoqijia, gnnnnng}@gmail.com ²Microsoft Research t-shtopl@microsoft.com

³Aroki Systems tarik@aroki.com ⁴National University of Singapore {prateeks, liangzk}@comp.nus.edu.sg

Abstract

Peer-to-peer (P2P) systems such as BitTorrent and Bitcoin are susceptible to serious attacks from byzantine nodes that join as peers. Research has explored many adversarial models with additional assumptions, ranging from mild (such as pre-established PKI) to strong (such as the existence of common random coins). One such widely-studied model is the *general-omission* model, which yields simple protocols with good efficiency, but has been considered impractical or unrealizable since it artificially limits the adversary only to omitting messages.

In this work, we study the setting of a synchronous network wherein peer nodes have CPUs equipped with a recent trusted computing mechanism called Intel SGX. In this model, we observe that the byzantine adversary reduces to the adversary in the general-omission model. As a first result, we show that by leveraging SGX features, we eliminate any source of advantage for a byzantine adversary beyond that gained by omitting messages, making the general-omission model *realizable*. Second, we present new protocols that improve the communication complexity of two fundamental primitives — reliable broadcast and common random coins (or beacons) — in the synchronous setting, by utilizing SGX features. Our evaluation of 1000 nodes running on 40 DeterLab machines confirms theoretical efficiency claim.

1 Introduction

Peer-to-peer systems such as BitTorrent [2], Symform [14], CrashPlan [5], StorJ [13], Tor [15] and Bitcoin [1] are becoming popular among users due to ease of accessibility. In such P2P systems, online users can simply volunteer as peers (nodes) to join the network. However, this exact property allows adversarial or Sybil peers to be a part of the network and exhibit a *byzantine* (malicious) behavior. The presence of byzantine adversaries is a major security concern in P2P systems. For example, recently, researchers have demonstrated that in a popular cryptocurrency — Bitcoin — byzantine nodes can collude to eclipse or partition the honest nodes leading to double-spending and selfish mining attacks [50, 67]. Further, byzantine nodes in anonymous P2P networks can become the entry and exit nodes of an honest node’s commu-

nication circuit, by advertising high-bandwidth connections and high-uptimes falsely [19]. These byzantine entry / exit nodes can selectively deny service or severely weaken the core anonymity properties of such systems as Tor, Cashmere and Hydra-Onions [15, 25]. In addition, byzantine nodes in the network can selectively forge, divert, delay or drop messages to disrupt the protocol execution. Therefore, designing robust P2P protocols continues to be an important research problem due to the attacks possible in a byzantine setting.

Researchers have extensively worked in the byzantine model to design solutions for fundamental P2P problems such as reliable broadcast and agreement among the peers [16, 17, 22, 24, 44, 45, 56, 71]. In a quest for efficient protocols that tolerate a larger fraction of malicious nodes, several failure models have been proposed which limit the capabilities of the byzantine adversaries. For instance, one such model is the *general-omission* model where the byzantine node can only omit messages that are either sent or received by it during the execution of a protocol [69, 72]. In this weaker adversarial model, designs with relatively simple and efficient protocols for reliable broadcast tolerating $\frac{N}{2}$ adversarial nodes are known [35, 49, 69, 72]. However, many of these models make strong assumptions, which are not always realistic and have not had a concrete basis for implementation.

Our approach. To this end, we study the possibility of using recent hardware-root-of-trust mechanisms for making previous adversarial models realizable in practical systems. We observe that emerging hardware, specifically Intel SGX, provides stronger trusted computing capabilities, which allow running hardware-attested user-level enclaves on commodity OSes [7–9, 39]. Enclaves provide hardware-isolated execution environment which guarantees that an application executing in an enclave is tamper-resistant and can be attested remotely. Assuming that SGX-like capabilities become commodity and widescale in end hosts, we ask if it is feasible to build robust P2P protocols. Our main observation is that by leveraging the capabilities of such a trusted hardware, one can restrict the behavior of byzantine adversaries to the *general-omission* model in synchronous networks [35, 49, 69, 72].

Specifically, we use four SGX features, i.e., enclave execution (F1), unbiased randomness (F2), remote attestation

(F3) and trusted elapsed time (F4). Based on these hardware features, we enforce six security properties (P1 - P6). First, we enforce execution integrity (P1), message integrity & authenticity (P2) and blind-box computation (P3) to restrict the attacker to not forge messages or deviate from the execution of the given protocol. Thus, the adversarial node can only delay, replay and omit messages. We further leverage lockstep execution (P5) and message freshness (P6) to reduce the adversarial model to the general-omission model, where byzantine nodes have no additional advantage than omitting to send / receive messages. In such model, P3 disallows the adversary to selectively omit messages based on the content. Lastly, the halt-on-divergence (P4) allows us to detect and eliminate peers that selectively omit messages based on identities of senders / receivers, thus in turn reducing round complexity and “sanitizing” the network. Leveraging these properties we can further achieve improvement for the efficiency of protocols. We present efficient designs for reliably broadcasting messages called *Enclaved Reliable Broadcast* (ERB) protocol and an unbiased common random generator called *Enclaved Random Number Generator* (ERNG) protocol. Both ERB and ERNG primitives can be used as building blocks to solve a wide range of problems in distributed systems, such as random beacons [74], voting schemes [65], random walks [48], shared key generation [46, 47], cryptocurrency protocols [61] and load balancing protocols [40, 75].

Results. Our work targets synchronous network where every machine is running an SGX-enabled CPU. Both of our protocols asymptotically reduce the round and communication complexity as compared to previous works in the byzantine model, and match with (or outperform) the results in general-omission model. For a network of size N , the round and communication complexity for ERB are $\min\{f + 2, t + 2\}$ and $O(N^2)$, where t / f ($f \leq t < \frac{N}{2}$) is the number of byzantine peers / peers actually behaving maliciously for one execution of ERB. The communication complexity of the basic ERNG is $O(N^3)$, and the optimized ERNG further reduces the complexity to $O(N \log N)$. We have implemented our solution and the source code is available online [10]. We evaluate both ERB and ERNG, and our experimental results match our theoretical claims.

Contributions. We summarize the main contributions of this paper as below:

- *Realizable General-Omission Model.* We leverage SGX features to reduce byzantine model to general-omission model, where byzantine nodes have no extra advantage than omitting messages.
- *Better Synchronous P2P Protocols.* By enforcing our properties, we can improve the efficiency of P2P protocols. As the first attempt, we propose efficient protocols for reliable broadcast (ERB) and unbiased random number generation (ERNG).
- *Security Analysis & Evaluation.* We provide security analysis and proof for our protocol constructions. Our experi-

mental evaluation confirms the theoretical expectations of our solutions.

2 Problem

Designing efficient solutions for P2P protocols in the byzantine setting is a widely-recognized problem with limited solutions [16, 17, 22, 44, 45, 71]. Our goal is to shed light on how SGX can aid to improve efficiency of synchronous P2P protocols. In this work, we take two fundamental problems as examples: 1) reliable broadcast and 2) common unbiased random number generator.

2.1 Problem Definition

In light of the previous works, we recall the standard definition of *reliable broadcast* [35, 69] and *common unbiased random number* [18] in the synchronous network:

Definition 2.1. (Reliable Broadcast). *A protocol for reliable broadcast in synchronous settings satisfies the following conditions:*

- *(Validity)* If the sender is honest and broadcasts a message m , then all honest nodes eventually accept m .
- *(Agreement)* If an honest node accepts m , then all honest nodes eventually accept m .
- *(Integrity)* For any message m , every honest node accepts m at most once, if m was previously broadcast by the sender.
- *(Termination)* Every honest node eventually accepts a message (m or \perp).

To define a *common unbiased random number generator*, we define the *bias* of any multi-variate function in a standard way [18].

Definition 2.2. (Unbiasedness). *Let $G : \{0, 1\}^{k \times N} \rightarrow \{0, 1\}^k$ be a deterministic multi-variate function that maps N elements in $\{0, 1\}^k$ to one element in $\{0, 1\}^k$. We define the bias of G , $\beta(G)$, as follows:*

$$\beta(G) = \max_{S \subseteq \{0, 1\}^k} \left(\max \left(\frac{E[S]}{E_G[S]}, \frac{E_G[S]}{E[S]} \right) \right),$$

where $E_G[S]$ is the expected number of values in $G(x_1, \dots, x_N) \in S$, and $E[S] = \frac{|S|}{2^k}$, which is the expected value when the output of G is distributed uniformly at random.

Definition 2.3. (Common Unbiased Random Number). *A protocol G generates a common unbiased random number r among N nodes if it satisfies the following conditions with high probability (w.h.p.):*

- *(Agreement)* At the end of the protocol, all the honest nodes agree on the same value r .
- *(Unbiasedness)* The bias of $\beta(G) = 1$.

For the analysis of protocols, we define the following complexities with respect to a single execution of the protocol.

- The *message / communication complexity* is defined as the total number of messages / bits transferred among all nodes in the worst case.
- The *round complexity* is defined as the number of executed rounds (or steps) in the worst-case.

2.2 Attacker Model

We consider a widely-studied standard synchronous model of P2P systems [16, 17, 22, 44, 45, 71]. In this model, our only new requirement is that every peer in the network uses an SGX-enabled CPU to run the P2P protocols. In a network of N nodes, the number of byzantine nodes t is strictly bounded under a fraction of $\frac{N}{2}$. The number of peers that actually behave maliciously for a particular execution of the protocol is $f(\leq t)$. Thus, a P2P network \mathcal{P} is composed of N peers $\mathcal{P} = \{p_1, \dots, p_N\}$ such that $N = 2t + 1$. Every peer p_i in the P2P overlay has an identifier id_i and can communicate with other peers using their ids. The underlying TCP/IP substrate is assumed to provide reliable message delivery within a known bounded delay say Δ . Moreover, we consider a round-based *synchronous* model where each round is equal to the time an honest node requires to send a message and receive a response. Every peer is directly connected to all other peers in the network and knows the network size N . To summarize, we assume: the network size is N (S1); the protocol starts synchronously (S2); the round time is 2Δ (S3); the number of byzantine nodes is limited upto $\frac{N}{2}$ (S4); the peers are connected to each other (S5). This is a prominently used model in the previous literature of distributed P2P systems [16–18, 48, 69, 72].

Our Model using SGX. In our model, a byzantine peer has a compromised or malware-ridden operating system but executes protocols using SGX enclaves [7, 8, 39]. Enclaves guarantee untampered execution in presence of malicious underlying software or co-processes. The byzantine nodes can take arbitrary software actions as long as it does not violate SGX guarantees.

Scope. Our focus is showing how to leverage SGX features to improve the efficiency of synchronous P2P protocols. Our model does not consider an adversary that can perform hardware attacks and break SGX security guarantees. We do not aim to prevent any information leakage through side-channels such as pagefaults, memory accesses or timing attacks to which SGX-enabled CPUs are known to be susceptible [57, 64, 80]. Indeed these problems are under investigation and recent research shows that defending against them is feasible. Existing solutions against these problems can directly apply to our work [62, 68, 76].

2.3 Strawman Solution & Attacks

Consider a strawman protocol for distributed random number generation using reliable broadcast, where the initiator broadcasts a random number m using an initialization message INIT to all the peers in a synchronous network (shown in Algorithm 1). If m is generated randomly and unbiasedly as well as reaches every honest node without being tampered, then all honest nodes will agree on the common unbiased random number m and the goal of the protocol is achieved. In Algorithm 1, upon receiving the INIT message, each peer

Algorithm 1: Strawman distributed random number generation protocol using reliable broadcast.

```

Input: A P2P network  $\mathcal{P}$  composed of  $N$  nodes, an initiator node  $id_{init}$ 
Output: A message  $\hat{m}$ 
1 Initialization:  $\hat{m} \leftarrow \perp$ ;  $S_m \leftarrow \emptyset$ ;  $rnd \leftarrow 1$ 
2 upon self_id is initiator:
3 get( $m$ ) //  $m$  is a random number
4  $\hat{m} \leftarrow m$ 
5 add self_id to  $S_m$ 
6 multicast INIT( $m$ ) to other peers
7 for  $rnd \leq t + 1$  do
8   upon receiving INIT( $m$ ):
9      $\hat{m} \leftarrow m$ 
10    add self_id and sender_id to  $S_m$ 
11    multicast ECHO( $m$ ) to other peers in round  $rnd + 1$ 
12   upon receiving ECHO( $m$ ):
13     if  $\hat{m} = \perp$  then
14        $\hat{m} \leftarrow m$ 
15       add self_id to  $S_m$ 
16       multicast ECHO( $m$ ) to other peers in round  $rnd + 1$ 
17     end
18     if  $m = \hat{m}$  and sender_id  $\notin S_m$  then
19       add sender_id to  $S_m$ 
20       if  $|S_m| = N - t$  then
21         accept  $\hat{m}$ 
22       end
23     end
24    $rnd \leftarrow rnd + 1$ 
25 end
26 if  $rnd > t + 1$  then
27   accept  $\perp$ 
28 end

```

further multicasts an ECHO message to all other peers. After receiving the ECHO messages from the majority of nodes, each peer accepts m as the final message \hat{m} . Note that if the initiator is honest, all honest nodes receive the message INIT during the first round and multicast ECHO messages at the beginning of the second round. In the second round, every honest node receives at least $N - t$ ECHO messages from $N - t$ honest nodes and maybe some byzantine nodes. Thus, after two rounds, every honest node will output the same value m from the initiator, which satisfies all the conditions of reliable broadcast in Definition 2.1. However, we show how a byzantine initiator and other byzantine peers can attack this protocol to violate Definitions 2.1 and 2.3.

Attacks by Byzantine Adversary. Byzantine initiator / peers can tamper with the execution of Algorithm 1 and forge the values of INIT and ECHO messages to perpetrate the following attacks.

A1 (Execution Deviation): For this attack, an adversary deviates from the control flow of the running program for the given protocol. The adversary can disregard essential conditions to jump to the desired instructions and execute them directly. For example, the adversary can skip all the conditions like Line 7 & 13 to directly multicast its ECHO value to parts of honest nodes but not all of them, to introduce equivocation to their final decisions. Moreover, the adversary can also repeat particular instructions to obtain an output she wants. For instance, if m is generated from a random source without being tampered during the execution of the protocol, an unbiased common random number can be agreed among all the peers in the network. A byzantine peer, however, can repeat

the step that generates m (Line 3) from the random source until it returns a favorable random number. Hence, the output is biased as per Definition 2.3.

A2 (Message Forgery): Suppose that the adversary does not deviate from the execution of the given protocol, she can still alter the data flow (including input / output and intermediate states) of the program to forge messages. As per Definition 2.1, a reliable broadcast protocol requires that if one honest node accepts message m then all honest nodes accept m . The adversary can tamper with the INIT and ECHO messages to violate this agreement property of the protocol. A byzantine initiator colluding with other byzantine peers in the network can tamper with Line 6, 11 and 16 in the algorithm such that some honest nodes receive most ECHO messages with m' while others with m . This results in a fraction of honest nodes assigning \hat{m} with m' and accepting m' , while other honest nodes accept m as the final output, thereby causing inconsistency in the network.

A3 (Selective Omission): Assume that the adversary does not deviate from the control flow (i.e., the execution) of the given protocol or tamper with the data flow to forge messages, she can still omit, delay and replay messages in this restricted model. For an omission attack, it has two types: one is based on the content of the transmitted message and the other is dependent on the identity of the sender / receiver. For the first type, the adversary can observe its generated or received random number m and selectively decide to drop or forward it to other nodes based on its value, which introduces a bias in the final output for the honest nodes. For example, if the adversarial peers receive or initiate a message m , which is not the favorable one, they can omit to relay the message to the other nodes, thus all honest nodes may finally agree on \perp instead of m . Further, to violate the agreement condition in Definition 2.1 and 2.3, the adversary can selectively decide to omit the message m depending on whether the destination peer is honest or malicious. It can broadcast m correctly to a few honest nodes and not send the message to the others for the last round. The honest nodes receiving m can multicast m to the others, but the others will not accept it as the execution ends. Thus, the honest nodes that do not receive a message will agree on \perp while others will agree on m .

A4 (Message Delay): Alternatively, to generate an unbiased common random number, every peer can broadcast its random number to all other peers using Algorithm 1. All peers can then XOR the random numbers in the final set to generate the output. To bias this final output, a byzantine peer can intentionally hold its random number until it receives inputs from all other honest peers [18]. In this way, the adversary can “look ahead” in the protocol, calculate the final output and then decide whether to participate in the protocol by sending its random number. If the final random number already favors the adversary then it does not participate in the protocol, otherwise it sends its message to all the peers. Note that, for $t < \frac{N}{2}$, all the byzantine adversaries can collude to introduce




Figure 1: Each peer consists of two entities: an Enclave and an OS. The OS models the operating system and memory. The Enclave models the isolated memory and the secure execution of a program. The sender Enclave_s can send a message via a secure channel to the receiver Enclave_r. The grey areas are secure against malicious OSes of byzantine nodes.

an exponential bias in the final value.

A5 (Message Replay): In the restricted model, the adversarial node can use a message m_{prev} from an instance of the protocol running in parallel, or which was run in the past to one (or more) honest node(s) and forward the correct message m to other honest nodes [59]. This results in an inconsistency where few honest nodes agree on m_{prev} and others agree on m , thereby violating the agreement condition.

3 Solution Overview

In this section, we put forward ideas using SGX features to enforce six security properties to restrict the capabilities (A1 - A5) of a byzantine adversary, as shown in Section 2.

3.1 SGX Features and Security Properties

We first start by recalling Intel SGX features (supported in both simulation and hardware modes in the latest version [7, 9]), which can also be provided by other trusted hardware.

F1: Enclaved Execution - SGX supports hardware-isolated memory region called enclaves such that a compromised underlying OS cannot tamper the execution of the code running inside this enclave.

F2: Unbiased Randomness - SGX provides a function `sgx_read_rand` that executes the `RDRAND` instruction to generate hardware-assisted unbiased random numbers.

F3: Remote Attestation - SGX allows a remote party to verify that an application is running in an enclave on an SGX-enabled CPU.

F4: Trusted Elapsed Time - SGX provides a function `sgx_get_trusted_time` that returns a trusted elapsed time in seconds relative to a reference point.

Abstractly, a peer can be considered as the composition of two entities: an OS and an Enclave as shown in Figure 1. The OS models the untrusted entity including the operating system and memory. It has access to all the system resources such as file system and network. The OS can arbitrarily invoke an enclave program and start its execution. The Enclave models the isolated memory space that loads the program and executes it securely. Thus, Enclave corresponds to the trusted entity of a peer. We illustrate how to enforce P1 - P6 properties using SGX features to thwart A1 - A5 attacks.

P1 (Execution Integrity): With remote attestation (F3), an enclave in one peer can verify the correctness of the running

program for the given protocol on the other nodes and whether it is executing on a valid SGX-enabled CPU or not. Moreover, F1 ensures that the execution in an enclave cannot be tampered with by the OS. F1 and F3 together enforce the execution integrity against A1. Hence, an adversary cannot deviate from the execution of the protocol in an enclave arbitrarily by skipping / repeating instructions to violate the control flow of the running program.

P2 (Message Integrity & Authenticity): In designing our protocols, we first perform a setup phase where each peer connects to every other in the network and then performs a series of steps. Analogous to P1, every enclave first uses F3 to verify the correctness of the protocol executing on other peers. Next, they generate public / private key pairs inside the enclaves and exchange the public keys with each other. Then all the messages transmitted between any two enclaves can be signed to ensure the integrity and authenticity against A2. Moreover, the internal states of the program are also protected using F1. Therefore, the integrity of all messages including input / output / intermediate states is guaranteed. In this case, it is clear that an adversary cannot forge valid messages to bias the honest nodes to make inconsistent decisions.

P3 (Blind-box Computation): F1 ensures that all intermediate states of the protocol's computation are hidden from the OS. Leveraging F2, the provided randomness is also hidden from the OS. This guarantees that the input state is hidden along with the intermediate states of the protocol's execution. We say in this case that the computation is a *blind-box computation*. As the adversarial node does not know the random number and given that the output of the computation is encrypted between the Enclave and the OS, she cannot selectively omit or drop messages based on their contents. Note that an important part of instantiating such a blind-box computation is the ability to instantiate a secure channel between two or more enclaves. In fact, enclaves can agree on a shared key to establish a secure channel using Diffie-Hellman key exchange. Nodes can then encrypt all the messages (including program's intermediate input / output) transmitted between each other to provide confidentiality against malicious OSes. Note that, establishing such a shared key in the enclaved setting is slightly weaker than the standard byzantine model, as the malicious operating system cannot access the shared secret keys and decrypt the exchanged messages due to F1. With P1 - P3, we can reduce the byzantine model to a restricted model, where an adversarial node can only replay, omit and delay messages.

P4 (Halt-on-Divergence): To mitigate selective omission based on nodes' identities (A3), we enforce a security mechanism called halt-on-divergence. This property halts any malicious node deviating from the protocol under some given condition. As an instance, if an adversarial node sends a message, but does not receive adequate responses, it will be forced to leave the current protocol execution. Halt-on-divergence mechanism should be incorporated through a specific ac-

knowledgment protocol instantiation in such a way that every malicious node will be forced to leave if the acknowledgment is not verified. In particular, we introduce an acknowledgment scheme where every receiver acknowledges the sender on receiving every valid message. A message sent over a secure channel is considered valid only if it contains the expected sequence and round number. Naturally, an acknowledgment is not sent for a replayed, omitted or delayed message. Since all honest receivers will reply with acknowledgment (ACK) messages on receiving valid messages, an honest sender should at least receive $t + 1$ ACK messages. Any node receiving less than $t + 1$ ACK messages will halt and leave the network.

The key idea here is to penalize any deviating adversary by churning the node out of the network. This effectively "sanitizes" the network. Thus, to remain a part of the network, every peer should send valid messages to the majority of the network. This property also aids honest nodes in the protocol to decide the final output early and finish the execution immediately.

P5 (Lockstep Execution): F4 allows us to realize a synchronized network across all rounds of a protocol. Each peer uses F4 to decide the correct value of the ongoing round and inserts this round number in all the sent messages. To detect *delay* attacks (A4), a peer simply matches the round number present in an incoming message with the current round number. This defense is hard in the byzantine model with public-key infrastructure even if it supports F1, since the OS can tamper with the relative time to either increase or decrease the rounds of a node. Therefore, having access to a trusted elapsed time functionality allows to perform lockstep execution and detect delay attacks in the restricted model.

P6 (Message Freshness): Similar to [59], we use sequence numbers to ensure message freshness and therefore defend against replay attacks (A5). The main challenge lies in ensuring secure exchange of the initial sequence numbers for each peer and ensuring that the sequence number remains untampered with during the entire intermediate states of the protocol execution. Using the secure channel, the peers securely exchange a nonce or a *sequence number*, which is incremented sequentially by the peer. The nonce is generated using F2 supported by SGX. This prevents the malicious adversary from tampering the initial nonce value to its own advantage. Note that the keys and initial sequence numbers exchange occur only once during the setup phase. If an adversarial node restarts or relaunches its enclave, all the data in the enclave will be removed. Since the enclave does not have the valid sequence number and round number, it cannot re-join the same or any on-going execution, which is equivalent to be considered as a new node for the protocol.

3.2 Overview of Our Results

In this work, we achieve the following results.

R1: *By enforcing (P1 - P6), we reduce the byzantine model to the general-omission model.*

| Protocol | Attacker Model | Network Size | Round Complexity | Comm. Complexity |
|-------------|----------------|--------------|------------------------|------------------------|
| PT [72] | Omission | $t + 1$ | $\min\{f + 2, t + 1\}$ | $O(N^3)$ |
| PR [69] | | $2t + 1$ | $\min\{f + 2, t + 1\}$ | $O(N^2)$ |
| CT [35] | | | $2t + 1$ | $O(N^2)$ |
| PSL [71] | Byzantine | $3t + 1$ | $t + 1$ | $O(\exp(N))$ |
| BGP [24] | | | $\min\{f + 2, t + 1\}$ | |
| BG [22] | | $4t + 1$ | $t + 1$ | |
| GM [44, 45] | | $3t + 1$ | $\min\{f + 5, t + 1\}$ | $O(\text{poly}(N))$ |
| AD15 [16] | Byzantine | $2t + 1$ | $\min\{f + 2, t + 1\}$ | $O(N^4)$ |
| AD14 [17] | | $2t + 1$ | $3t + 4$ | $O(N^4)$ |
| ERB | | Byz. + SGX | $2t + 1$ | $\min\{f + 2, t + 2\}$ |

Table 1: Round complexity and communication complexity for reliable broadcast in synchronous network.

| Protocol | Network Size | Round Complexity | Comm. Complexity |
|----------------|--------------|------------------|------------------|
| AS [18] | $6t + 1$ | $O(N)$ | $O(N^3)$ |
| AD14 [17] | $2t + 1$ | $O(N)$ | $O(N^4)$ |
| Basic ERNG | $2t + 1$ | $O(N)$ | $O(N^3)$ |
| Optimized ERNG | $3t + 1$ | $O(\log N)$ | $O(N \log N)$ |

Table 2: Round / communication complexity for random number generation protocols in synchronous distributed systems.

By enforcing P1 - P3, we first reduce byzantine model to a restricted model, in which byzantine nodes can only delay / omit / replay messages. We defer the formalization and proof to Appendix A. We believe that the formalization, while based on traditional cryptographic primitives, provides a new conceptual framing of SGX-enabled CPUs security features, and may be of independent interest. By applying P5 and P6, we further confine the adversarial nodes into the general-omission model.

R2: We propose an efficient reliable broadcast protocol (ERB) with early stopping, which improves communication complexity from $O(N^3)$ to $O(N^2)$ (refer to Section 4).

For this result, we leverage four properties. First, P1 - P3 ensure that the adversarial nodes cannot forge messages and deviate from the execution of the protocol. Second, we leverage P4 to show that ERB can broadcast a message to the entire network in $\min\{f + 2, t + 2\}$ rounds with better performance as shown in Table 1.

R3: We propose a new unbiased random number generation protocol (ERNG) with communication complexity $O(N^3)$ for the basic version, or $O(N \log N)$ for the optimized one, as shown in Table 2 (refer to Section 5).

With P3 and P5, our unoptimized ERNG solution directly runs our ERB protocol as a sub-routine on the entire network to agree on a random number generated using F2. It has round and communication complexity of $O(N)$ and $O(N^3)$, respectively. We present an optimized version of ERNG by reducing the byzantine fraction from $\frac{N}{2}$ to $\frac{N}{3}$, and forming a cluster of peers within the network. Leveraging the trusted randomness F2 and P3, we can sample a small set of nodes forming a representative cluster. The ERB protocol is executed within this small cluster to generate the final unbiased random number. The round and communication complexity of this optimized ERNG is further reduced to $O(\log N)$ and $O(N \log N)$. Note that the optimized version of ERNG only applies when the size of the network is large enough.

4 Enclaved Reliable Broadcast Protocol

We propose an *enclaved reliable broadcast* (ERB) in the synchronous model using SGX features. The transmitted message, val , between any two peers has the format: $\text{val} := \langle \text{type}, \text{id}, \text{seq}, m, \text{rnd} \rangle$, where $\text{type} \in \{\text{INIT}, \text{ECHO}, \text{ACK}\}$ and rnd represents the current round of the ERB protocol. If $\text{type} = \text{INIT}$, then the initiator peer id_{init} is initiating the broadcast by sending the message m with sequence number seq_{init} at round rnd . If $\text{type} = \text{ECHO}$, it means that its sender knows that id_{init} has sent m , as it has already received either a value with INIT or ECHO for the first time. Finally, if $\text{type} = \text{ACK}$, it means that the peer acknowledges that it has already received either INIT or ECHO values from the sender. We introduce three functions Halt, Multicast and Wait:

- $\text{Halt}(\text{st})$: is a function that sets the state st to \perp .
- $\text{Multicast}(\text{id}_i, \text{val})$: is a functionality that multicasts the value val from the sender p_i to the receiver p_j , for all $j \in [N] \setminus \{i\}$.
- $\text{Wait}(\tau)$: is a function that has as an input the current elapsed time τ in the ongoing round, and suspends the protocol for $(2\Delta - \tau)$ seconds.

Note that Halt function enforces the *halt-on-divergence* property (P4) that we have introduced in Section 3. When the state of the node is set to \perp the node halts on-divergence and is ejected from the P2P network \mathcal{P} . For the sake of exposition, we write $\text{Wait}(\text{rnd})$ in the code description, we say in this case that the protocol waits until the end of the round rnd .

4.1 ERB details

Prior to running the very first instance of the ERB protocol, there is a setup phase. The setup is performed whenever the program (ERB) needs to be updated or changed. We detail the setup phase followed by the explanation of our algorithm.

Setup Phase: Every pair of sender and receiver peer use remote attestation (F3) along with enclaved execution (F1) to verify the correctness of the execution, and therefore enforcing P1 - P3. Then they establish a secure channel using Diffie-Hellman key exchange. This setup enforces P1 - P3, which restricts the byzantine nodes to only omit, replay and delay messages. Next, each peer picks at random a sequence number such that $\text{seq}_s, \text{seq}_r \stackrel{\$}{\leftarrow} \{0, 1\}^k$ and send it to each other. That is, every node has to store the sequence numbers of all other nodes in \mathcal{P} . Finally, every node sets the variable rnd to the value 1. The overhead of the setup is in $O(N^2)$ while the storage overhead per node is in $O(N)$.

Initialization Phase: An initiator node first multicasts the value $\text{val} = \langle \text{INIT}, \text{id}_{\text{init}}, \text{seq}_{\text{init}}, m, \text{rnd} \rangle$, where seq_{init} is the sequence number of the initiator node, and rnd is the round number. The round rnd is first initialized to 1, the enclave will now increment the rnd after every 2Δ seconds—we take advantage of the elapsed time feature of SGX to tie a round to an interval of 2Δ seconds.

Echo Phase: Until round $t + 2$, if a node receives an INIT or ECHO message for the first time, it performs the following actions: (1) start the local clock and initialize the round rnd to 1, the round will increment every 2Δ seconds, (2) if both rnd and seq are consistent with the expected values, it will store the message m , else it just ignores it and treats it as an omitted message. If there is no delay or replay detected, then it multicasts an ECHO message to all nodes at the end of the current round. If the node has already received a valid ECHO message from a distinct node, it will only add the sender's identifier into the set S_{echo} . Recall that at the end of the setup phase, all honest nodes have the same copy of the sequence number of all honest nodes. After every valid instance of the protocol, nodes will increase all sequence numbers by 1.

Decision Phase: If the node has received at least $t + 1$ correct ECHO messages from distinct nodes, i.e., $|S_{\text{echo}}| = t + 1$, then the node accepts \hat{m} . After $t + 2$ rounds, if the node has not received adequate distinct ECHO messages, it accepts $\hat{m} := \perp$. Every multicast requires the node to receive at least $t + 1$ ACK messages, else the node churns out itself using Halt.

4.2 Analysis

In Algorithm 2, if a byzantine sender decides to omit a message, it will not receive a corresponding ACK message as the sent messages never reach the receiver peer. The sender Enclave_s detects that the underlying OS_s is byzantine if it does not receive at least $t + 1$ ACK messages. On failing to receive majority ACK messages, Enclave_s executes the Halt function as per our algorithm and churns itself out of the network based on our halt-on-divergence property (P4). By leveraging the P4, any node can actively detect its own anomalous behavior instead of relying on other nodes to send messages every round to passively identify the anomaly. This results in communication complexity for anomaly detection decreased from $O(N^2)$ to $O(N)$ and the overall complexity is reduced to $O(N^2)$, compared to previous passive-detection approaches, e.g., Perry *et al.*'s work [72]. Here we state our main theorem below and defer the detailed proof to Appendix B.

Theorem 4.1. *If $N \geq 2t + 1$, ERB is a reliable broadcast protocol as defined in Definition 2.1.*

ERB Performance Analysis. Algorithm 2 has a worst-case round complexity equal to $t + 2$ with communication complexity in $O(N^2)$ and $t < \frac{N}{2}$ byzantine nodes. This only occurs if the byzantine peers delay the instance for t rounds before sending the message to at least one honest node. However, in this case, the round complexity is equal to $f + 2$ rather than $t + 2$ as the delay is only in function of the number of byzantine nodes f . On the other hand, byzantine nodes can also decide to not send the message to any honest node, and then the round complexity is $t + 2$ with $O(t)$ communication complexity.

Algorithm 2: ERB: Enclaved reliable broadcast protocol (for a node id_i with the initiator id_{init} sending a message m and a sequence number seq_{init}).

Input: A P2P network \mathcal{P} composed N nodes, a message m and a sequence number seq_{init} for the initiator id_{init}

Output: A message \hat{m}

```

• initialization:  $\hat{m} \leftarrow \perp$ ;  $S_{\text{echo}} \leftarrow \emptyset$ ;  $\text{rnd} \leftarrow 1$ 
• upon  $\text{id}_i = \text{id}_{\text{init}}$  and  $\text{st}_i \neq \perp$ :
   $\hat{m} \leftarrow m$ ;
   $S_{\text{echo}} \leftarrow S_{\text{echo}} \cup \{\text{id}_{\text{init}}\}$ ;
  Multicast( $\text{id}_{\text{init}}, (\text{INIT}, \text{id}_{\text{init}}, \text{seq}_{\text{init}}, m, \text{rnd})$ );
• for  $\text{rnd} \leq t + 2$  do
  • upon receiving  $\langle \text{INIT}, \text{id}_{\text{init}}, \text{seq}, m, \text{rnd}' \rangle$  from  $\text{id}_{\text{init}}$ :
    if  $\text{rnd}' = \text{rnd}$  and  $\text{seq} = \text{seq}_{\text{init}}$  then
      send  $\langle \text{ACK}, \text{id}_{\text{init}}, \text{seq}, H(m), \text{rnd} \rangle$  to  $\text{id}_{\text{init}}$ ;
       $\hat{m} \leftarrow m$ ;
       $S_{\text{echo}} \leftarrow S_{\text{echo}} \cup \{\text{id}_{\text{init}}\} \cup \{\text{id}_i\}$ ;
      Wait( $\text{rnd}$ ) then Multicast( $\text{id}_i, \langle \text{ECHO}, \text{id}_{\text{init}}, \text{seq}, m, \text{rnd} + 1 \rangle$ );
    end
  • upon receiving  $\langle \text{ECHO}, \text{id}_{\text{init}}, \text{seq}, m, \text{rnd}' \rangle$  from peer  $\text{id}_j$ :
    if  $\text{rnd}' = \text{rnd}$  and  $\text{seq} = \text{seq}_{\text{init}}$  then
      send  $\langle \text{ACK}, \text{id}_{\text{init}}, \text{seq}, H(\text{val}), \text{rnd} \rangle$ , where
       $\text{val} = \langle \text{ECHO}, \text{id}_{\text{init}}, \text{seq}, m, \text{rnd} \rangle$  to peer  $\text{id}_j$ ;
      if  $\hat{m} = \perp$  then
         $\hat{m} \leftarrow m$ ;
         $S_{\text{echo}} \leftarrow S_{\text{echo}} \cup \{\text{id}_j\}$ ;
        Wait( $\text{rnd}$ ) then
          Multicast( $\text{id}_i, \langle \text{ECHO}, \text{id}_{\text{init}}, \text{seq}, m, \text{rnd} + 1 \rangle$ );
        end
      if  $\text{id}_j \notin S_{\text{echo}}$  then
         $S_{\text{echo}} \leftarrow S_{\text{echo}} \cup \{\text{id}_j\}$ 
        if  $|S_{\text{echo}}| = N - t$  then
          accept  $\hat{m}$ ;
        end
      end
    end
  • upon Multicast( $\text{id}_i, \text{val}$ ):
    send  $\text{val}$  to  $\text{id}_k$ , for all  $k \in [N] \setminus \{i\}$ ;
    receive  $N_{\text{ack}}$  acknowledgements  $\langle \text{ACK}, \text{id}_{\text{init}}, \text{seq}, H(\text{val}), \text{rnd}' \rangle$ , where
     $\text{rnd}' = \text{rnd}$  and  $\text{seq} = \text{seq}_{\text{init}}$ ;
    if  $N_{\text{ack}} < t$  then
      Halt( $\text{st}_i$ );
    end
  •  $\text{rnd} \leftarrow \text{rnd} + 1$ ;
end
• if  $|S_{\text{echo}}| < N - t$  then
   $\hat{m} \leftarrow \perp$ ;
  accept  $\hat{m}$ ;
end
•  $\text{seq}_{\text{init}} \leftarrow \text{seq} + 1$ ;

```

5 Enclaved Random Number Generation

We present our algorithm that generates an unbiased common random number called *enclaved random number generation* (ERNG).

5.1 Unoptimized ERNG

We detail our unoptimized ERNG in Algorithm 3. At a higher level, every node generates a random number from the enclave, and then performs ERB protocol to broadcast to every node. According to Theorem B.1, all honest nodes in this case will receive the random numbers from all honest nodes after $t + 2$ rounds, and may eventually receive several random numbers from other byzantine nodes. According to the validity requirement, for each ERB instance, every honest node will accept a random number from its initiator or \perp so that all honest nodes have the same final set S_{final} of random numbers. By performing exclusive disjunction (or XOR) of

Algorithm 3: Unoptimized-ERNG: Unoptimized enclaved unbiased random number generation protocol executed by peer p_i .

Input: A P2P network \mathcal{P} composed of N nodes
Output: A unbiased random number r

- initialization: $S_{\text{final}} \leftarrow \emptyset; \text{rnd} \leftarrow 1$
- for** $\text{rnd} \leq t + 2$ **do**
- **if** $\text{rnd} = 1$ **then**
- initiate ERB with inputs $m_i \xleftarrow{\$} \{0, 1\}^k$ and seq_i ;
- end**
- if** $2 \leq \text{rnd} \leq t + 2$ **then**
- execute ERB instances and wait for the output
 ($M_i = \{m_1, \dots, m_i\}$);
- end**
- $\text{rnd} \leftarrow \text{rnd} + 1$;
- end**
- $S_{\text{final}} \leftarrow M_i$;
- $\text{seq}_j \leftarrow \text{seq}_j + 1$, for all $j \in [N]$
- accept $r = \bigoplus_{v \in S_{\text{final}}} v$.

all received random numbers, every honest node obtains an *unbiased common* random number eventually.

Unbiasedness and Randomness Analysis. We describe the main intuition behind the common unbiasedness and randomness of our ERNG’s output and defer formal details to Appendix C. To bias the random value, the adversary may perform several attacks. It can first try to directly forge the random number, however, this is restricted as per execution integrity (P1) and message integrity (P2) enforced by F1 and F3. An adversary can force the program to generate a local random number of its choice. However, each enclave generates an unbiased random number from SGX-enabled CPU instruction `RDRAND` using F2. It is not possible to bias the source of randomness based on the hardware guarantees. Our blind-box computation (P3) together with the secure channel guarantee that an adversary cannot selectively omit its random number based on its value with the goal to bias the output.

One adversarial strategy is to learn the final output and then decide whether to participate or not in the protocol, as in Attack A4. From Algorithm 3, all honest nodes output the final value after round $t + 2$. In order to bias the final value, the adversary should perform the following steps within round number $t + 2$: (1) learn the XOR of random numbers from honest nodes, (2) decide whether to participate or not based on the final value, (3) and multicast its number to honest nodes. In Algorithm 3, the final XOR operation executes only when $\text{rnd} > t + 2$. The execution integrity (P1) ensures sequential execution of our protocol. This property restricts the adversary from directly jumping to the step that computes the XOR operation and learn the result before other honest nodes generate the final output. Next, the lockstep execution (P5) enforced by the elapsed time feature (F4) allows us to bound the time for each round, even on a byzantine peer. Therefore, the adversary cannot look ahead and compute the final output before the last round. If the adversary decides to delay its own random number based on the computed final value, the adversarial random number will be neglected by all honest peers as it will reach after $t + 2$ round. Combining P1, P5 and P3, it is not possible for the byzantine adversary to achieve

steps (1) and (3) simultaneously.

For clarity and without any loss of generality, we model Algorithm 3 as a multi-variate function $G : \{0, 1\}^{k \times N} \rightarrow \{0, 1\}^k$ that maps N elements in $\{0, 1\}^k$ to one element in $\{0, 1\}^k$ such that $G(x_1, \dots, x_N) = \bigoplus_{i=1}^N x_i$.

Theorem 5.1. *The bias of G $\beta(G) = 1$.*

We defer the proof to Appendix C.

5.2 Optimized ERNG

Next, we illustrate the main steps behind our optimized ERNG and defer the pseudo-code details to Appendix D. In this section, we consider that at most $t \leq \frac{N}{3}$ nodes of the network can be byzantine. In this case, ERNG terminates after $\gamma + 4$ rounds, where γ is a statistical parameter. The intuition behind our optimization can be formulated as follows: we first notice that if we select uniformly at random a subset of nodes from \mathcal{P} , we can still guarantee w.h.p. the existence of an honest majority within this smaller representative cluster. By leveraging F2 to generate a random number and blind-box computation (P3), we can sample a set of peers forming a representative cluster. The main remaining question, therefore, is how large this cluster should be. As a starting point, note that if the cluster size is equal to $\frac{2N}{3}$, the probability of having an honest majority is equal to one. This already suggests that the cluster size can be chosen to be smaller. Conceptually, the protocol can be decomposed into three main steps:

Cluster Selection: The purpose of this step is to construct a representative cluster of the entire P2P network. The cluster will consist of nodes selected uniformly at random from \mathcal{P} . At round 1, every node picks uniformly at random a number from $\{0, \dots, \frac{N}{2\gamma} - 1\}$ using SGX (F2). This operation is protected leveraging property P3 in such a way that the computation is hidden from the OS. If the random number equals 0, then the node is *chosen* to be part of the cluster, and then it multicasts a CHOSEN message to all nodes in \mathcal{P} . Upon receiving the CHOSEN message, every chosen node adds the identifier of the sender to its own set S_{chosen} . The size of the set S_{chosen} represents the size of the cluster.

ERB Instances: We first detail a pseudo-solution and then detail our main construction in Algorithm 4 in Appendix D. In round 2, the nodes constituting the cluster will each generate a random number and broadcast it *only* to the nodes constituting the cluster (i.e., peers’ identifiers in S_{chosen}). That is, every node in the cluster will run an independent ERB instance. The intuition behind these multiple instances is the following: for the broadcast to be effective, at least one broadcast instance has to succeed in that the accepted message is different from \perp . However, the complexity of such solution is cubic in $O(|S_{\text{chosen}}|^3)$ which can be a handicap in term of efficiency. As a solution, we incorporate a two-phases clustering. The idea behind this choice is the following: in order to generate a random number we only require one honest node to output a random number r (otherwise the ERNG protocol may out-

put \perp). We can then proceed to select just a few number of nodes to perform the ERB protocol. As long as at least one of these nodes is honest, the correctness of our ERNG holds. Concretely, to generate the second representative cluster, we perform the following: from nodes in S_{chosen} , we uniformly pick at random a value from $\{0, \dots, \gamma' - 1\}$, where γ' is a parameter in function of γ that verifies $\gamma' \leq \gamma$. The peers that output a random number equal to zero will be the only peers able to initiate the ERB protocol. We will show that this strategy will greatly decrease the communication complexity and defer its analysis to Appendix D. Note that this phase lasts for $\gamma + 2$ rounds when all ERB instances terminate.

Selection Decision: At the end of the broadcast phase, the node of the clusters will have each a set containing eventually several random numbers. Note that, as ERB is a reliable broadcast primitive, we know that all honest peers in the cluster will have the same set of random numbers. Once a node in \mathcal{P} receives at least $\gamma + 1$ sets of random numbers, M_K , originating from the nodes in the cluster, it will output the set M_K as S_{final} . All honest nodes will output the same set under the assumption that there is a majority of honest nodes in the cluster. Finally, the random number equals the XOR value of all random numbers in S_{final} .

5.3 Analysis

We present the proofs for the Lemma and Theorems below in Appendix D.

Lemma 5.1. *If up to $t = \frac{N}{3}$ nodes are byzantine, then with at least $1 - \text{negl}(\gamma)$ probability, the representative cluster has more than γ honest nodes, and less than γ byzantine nodes.*

Theorem 5.2. Agreement: *All honest nodes eventually agree on the same common set S_{final} in ERNG.*

Theorem 5.3. Unbiasedness: *The output of the ERNG protocol is an unbiased random number.*

ERNG Performance Analysis. Note that in ERNG, $O(\gamma)$ nodes will be chosen to form the first representative cluster and therefore run $O(\gamma)$ Multicast functions. The communication complexity of this first step is $O(\gamma^2)$. Then, among this first representative cluster, a second cluster will be composed such that all nodes of this cluster will run each an ERB instance. If the size of the second representative cluster is $O(\sqrt{\gamma})$ (as shown in Corollary D.1 in Appendix D), then the communication complexity of this step is $O(\gamma^2 \cdot \sqrt{\gamma})$. Finally, the member of the first representative cluster will multicast the output of the ERB instances to all peers in \mathcal{P} . The communication complexity of this final step is $O(N \cdot \gamma)$. That is, overall, the communication complexity of ERNG equals $O(N \cdot \gamma + \gamma^{\frac{5}{2}})$. Based on Lemmas D.1 and D.2, if N is large such that it verifies $\gamma \in o(N)$, then we can set $\gamma \in O(\log N)$. In this case, the communication complexity and round complexity of ERNG are equal to $O(N \log N)$ and $O(\log N)$.

6 Evaluation

Implementation. We have implemented a prototype of ERB, unoptimized ERNG and ERNG in C/C++ using Intel SGX's Linux SDK [8]. The implementation contains 4030 lines of code (LOC) measured using CLOC tool [4]. Our prototype implementation is open source and available online [10]. We re-use the ported OpenSSL library including cryptographic utilities (`libcrypto` available with Intel SDK), to perform Diffie-Hellman key exchange and AES encryption/decryption. We use `boost` [3] library to implement the communications between any two nodes and use Google `protobuf` libraries [11] and `rapidjson` [12] to serialize data.

Experimental Setup. We use the DeterLab network testbed for our experiments [6]. It consists of 40 servers running Ubuntu 14.04 with dual Intel(R) Xeon(R) hexacore processors running at 2.2 GHZ with 24 cores and 24 GB of RAM. All machines are connected and share the same link with the bandwidth of 128MBps. Every node in our protocol takes up to 1 - 800 MB memory which limits the maximum number of nodes to 2^{10} in our experiments. Due to the limited number of machines in our testbed, we have to run multiple nodes on each machine, thus we use SGX simulation mode¹ for our program and use a simulated Intel attestation service (IAS).


Evaluation Methodology. To evaluate the correctness of our protocols, we measure the round complexity (time to terminate) and communication complexity (network traffic) for ERB, unoptimized-ERNG and ERNG, by varying the number of nodes from 2^2 to 2^{10} . We have highly optimized our system to handle dynamic ports allocations to handle a larger number of nodes within one machine (order of 25 nodes per machine). Part of our results reported in this section are for the *optimistic* case where all nodes behave honestly. We evaluate the round complexity of ERB while varying the number of byzantine nodes in the network up to $\frac{1}{4}$ of the entire network composed of 512 nodes. We also compare our experiment results for the traffic size with theoretical ones to verify if they match our asymptotic analysis.

6.1 ERB Evaluation


Honest Termination: Constant Scalability. Determining the termination of ERB is essential to validate our reliable broadcast primitive. Fig. (2a) shows that the termination time, in the case of an honest initiator, is nearly equal to twice the value of one round. This validates our theoretical results where we show that ERB finishes in 2 rounds when the initiator is honest. The small increase at 2^8 is purely due to the bandwidth bottleneck of our testbed, as the nodes share the same link.

Traffic Size: Quadratic Scalability. Fig. (3a) demonstrates that the communication complexity quadratically increases in function of the number of peers in \mathcal{P} (note that the x-axis is


¹ All SGX features we use are supported in the simulation mode and F4 is supported in seconds.



(a) Termination of ERB slightly increase with the number of peers.




(b) Termination time of ERNG in function of the number of nodes in \mathcal{P} .




(c) Time termination of ERB linearly increase with the number of byzantine nodes in \mathcal{P} .


Figure 2: Termination time in seconds for ERB, both unoptimized and optimized versions of ERNG in honest and byzantine network with different fractions.



(a) Communication of ERB in function of the number of nodes in \mathcal{P} .



(b) Communication of ERNG in function of the number of nodes in \mathcal{P} .



(c) Communication of ERB in function of different byzantine peers in \mathcal{P} .

Figure 3: (Th) theoretical and (Ex) experimental comparisons off network overall communication bandwidth in MB for ERB, both unoptimized (ERNG-0) and optimized (ERNG-1) versions of ERNG in honest and byzantine network with different fractions.

logarithmic). The message size of INIT and ACK is around 100 Bytes and 80 Bytes, respectively. For 1024 nodes in \mathcal{P} , the traffic size equals 277 MB. We show that this result matches our theoretical expectation.

6.2 ERNG Evaluation

Honest Termination: Limited Scalability. We show in Fig (2b) that ERNG termination remains slightly constant from 2^2 to 2^7 and then increases afterwards. Unfortunately, this does not reflect our theoretical findings and this is mainly due to the limitation of our testbed, namely, the upper bound on the communication link of 128MBps that all nodes have to share. For small values of peers N , the communication complexity of the unoptimized ERNG is cubic in N , while the optimized version is also (nearly) cubic for smaller values of N . Given a fixed bandwidth, this explains why the termination increases for larger values of N to reach 103 s for one instance.

Traffic Size: Cubic Scalability. Fig. (3b) demonstrates that the communication complexity cubically increases in function of the number of peers in \mathcal{P} for the unoptimized ERNG. Our theoretical results back up our experimental result. For ERNG as the bandwidth links get overflowed much faster, we limited our experiments to 512 nodes. For the optimized ERNG, small values of the number of peers in the network did not allow us to optimally select a cluster size that can guarantee w.h.p. the agreement. In this case, we fix the cluster to be $\frac{2}{3}$ of the network and we show that the traffic size decreases and has a 60% improvement over the unoptimized one. Note that

this result can get much better for a larger number of peers in realistic settings. Here, we draw our theoretical curve for the ideal evaluation which can be guaranteed only for larger N .

6.3 Byzantine case

In Fig (2c), we show that the termination time of ERB linearly increases with the number of byzantine nodes behaving maliciously in the current instance. We gradually increase the fraction of byzantine nodes from $\frac{1}{512}$ to $\frac{1}{4}$. As a strategy of byzantine nodes, we have taken into consideration the worst-case where byzantine nodes create a chain (a byzantine sends its message to only one byzantine node each round and then gets eliminated) in order to delay the termination as much as possible. In the case of $\frac{1}{4}$ byzantine fraction, the ERB termination takes 389 seconds while it only takes 4 seconds in the honest case. For traffic size, if the number of byzantine nodes increases, the communication complexity of ERB decreases as shown in Fig. (3c). This is mainly due to the halt-on-divergence property that will eject the nodes whenever it behaves maliciously. That is when an honest node multicasts a message, the eliminated byzantine node will not acknowledge this message which greatly reduces the communication complexity. For example, for $\frac{1}{4}$ byzantine fraction in a 512-node network, the traffic size equals 35 MB, while in an honest node instance, it is equal to 69 MB, a 50% decrease.

7 Related Work

Reliable broadcast has been extensively investigated in various adversarial models. In our work, we show how Intel SGX

improves the efficiency of existing protocols in these models, renewing interest in studying these protocols with SGX-based implementations.

Reliable Broadcast: Reliable broadcast has been extensively studied since the 1980s, and is closely related to the problem of byzantine agreement (BA). Several excellent surveys on the problem are available [55, 78]. Byzantine agreement can also achieve reliable broadcast [26, 29, 31, 52, 63, 66, 73, 78]. For the asynchronous network, Bracha’s classic reliable broadcast protocol requires $O(N^2)$ communication complexity and tolerates up to $\frac{N}{3}$ byzantine nodes [27, 28]. Cachin and Tessaro [32] leverage erasure codes to improve efficiency and reduce communication complexity. However, as the time is not bounded, messages may incur arbitrary delays, and most protocols do not guarantee terminating runs, except under some special assumptions such as sharing a “common coin” [26, 73].

Lamport *et al.* and Pease *et al.* propose protocols terminating within $t + 1$ rounds and tolerating up to $\frac{N}{3}$ byzantine nodes, but with exponential communication complexity [56, 71]. Berman *et al.* achieve $O(\text{poly}(N))$ communication complexity but only tolerating up to $\frac{N}{4}$ byzantine nodes [22]. Garay *et al.* later present a BA protocol terminating within $\min\{f + 5, t + 1\}$ rounds [44, 45].

To tolerate a larger fraction of byzantine nodes, additional assumptions are often needed. A common assumption is that of having a one-time trusted dealer that pre-deploys PKI in the infrastructure. This assumption, for instance, allows digital signatures to be used for *authentication*, wherein a message claimed to be sent by a node A can be assured to be originating from A [41, 43, 53, 56]. This weakens the capabilities of the byzantine adversary, which cannot forge messages on behalf of honest nodes. Researchers have proposed protocols to use digital signatures to boost the resilience from $\frac{N}{3}$ to $N - 1$, but the communication complexity is still large, i.e., $O(\exp(N))$ and $O(N^3)$ [41, 56]. Katz *et al.* extend the work of Feldman and Micali [42] to employ authenticated channels, and present protocols tolerating $\frac{N}{2}$ byzantine nodes with $O(\text{poly}(N))$ complexity [53]. Fitzi *et al.* also give an authenticated BA protocol that beats this bound ($\frac{N}{2}$) but under specific number-theoretic assumptions [43]. Abraham *et al.* provide a solution with early stopping ($\min\{f + 2, t + 1\}$) and polynomial complexity [16]. In this work, we use SGX features to reduce the byzantine model to the general omission model, and further propose ERB to achieve $\min\{f + 2, t + 2\}$ round complexity and $O(N^2)$ communication complexity.

Researchers also have proposed byzantine fault-tolerant algorithms using trusted services, such as by using trusted computing primitives, primarily focusing on making PBFT more efficient [20, 34, 36–38, 58, 60, 79]. These works have observed similar relation to crash-fault-tolerant protocols, as we have. For example, Chun *et al.* introduce an attested append-only memory (A2M) to remove the ability of adversarial replicas to equivocate without detection, which helps to increase the resilience from $\frac{N}{3}$ to $\frac{N}{2}$ [36]. However, these works have

concentrated on handling asynchronous protocols with weak time assumptions like PBFT. In this paper, in contrast to previous approaches, we work on the round-based synchronous model. Our work extends these ideas to detecting and remediating failures of synchronous network assumptions (e.g. our lockstep execution and halt-on-divergence). Additionally, we investigate the use of our blind-box execution primitive in our new distributed RNG protocol which is bias-resistant, and more efficient using secure sampling for cluster creation. We leave the extension of applying our properties and primitives to asynchronous protocols for future work.

Distributed RNG: Generating common coins in a distributed manner for randomized BA in asynchronous networks can also be used for generating unbiased random numbers [23, 30, 73]. However, these protocols either require a trusted dealer to set up the initial states of different nodes or pre-distribute data to the nodes in the network. Other works employing asynchronous verifiable secret sharing (AVSS) protocols do not have the trusted dealer, but can probabilistically execute with errors [21, 26, 33, 77]. Most of these works employ some cryptographic primitives that, in most case, can be considered heavy-weight and performance unfriendly. Awerbuch *et al.* propose a solution that tolerates up to $\frac{N}{6}$ byzantine nodes, with $O(N)$ round complexity and $O(N^3)$ communication complexity [18] to generate a random number with a constant bias. Other works, such as Andrychowicz *et al.*’s one, generate a common random number based on proof of work [17] with $O(N^4)$ communication complexity, but the output can eventually be biased. Moreover, the large communication cost for most of these approaches prevents scalability to a large number of nodes. We present more efficient (with $O(N \log N)$ communication complexity) and unbiased RNG generation for the synchronous network case.

8 Conclusion

The recent availability of Intel SGX in commodity laptops and servers provides a promising research direction for advancing the area of P2P systems. Our main observation is that leveraging SGX features can restrict a byzantine model to a general-omission model in synchronous systems. We highlight that using SGX we can improve the efficiency of P2P protocols such as reliable broadcast and unbiased random number generator in synchronous settings.

References

- [1] Bitcoin. <https://bitcoin.org/en/>, Accessed: 2020.
- [2] BitTorrent. <http://www.bittorrent.com/>, Accessed: 2020.
- [3] Boost C++ library. <http://www.boost.org/>, Accessed: 2020.

- [4] CLOC. <http://cloc.sourceforge.net/>, Accessed: 2020.
- [5] CrashPlan. <http://www.code42.com/crashplan/>, Accessed: 2020.
- [6] DeterLab. <https://www.isi.deterlab.net/index.php3>, Accessed: 2020.
- [7] Intel Software Guard Extensions. <https://software.intel.com/en-us/sgx>, Accessed: 2020.
- [8] Intel Software Guard Extensions for Linux OS. <https://01.org/intel-softwareguard-extensions>, Accessed: 2020.
- [9] Intel software guard extensions sdk for linux os. https://download.01.org/intel-sgx/linux-1.9/docs/Intel_SGX_SDK_Developer_Reference_Linux_1.9_Open_Source.pdf, Accessed: 2020.
- [10] P2P using SGX. <https://bitbucket.org/P2PUsingSGX/p2pusingsgx>, Accessed: 2020.
- [11] Protocol Buffers - Google's data interchange format. <https://github.com/google/protobuf>, Accessed: 2020.
- [12] RapidJSON. <http://rapidjson.org/>, Accessed: 2020.
- [13] Storj.io. <http://storj.io/>, Accessed: 2020.
- [14] Symform. <http://www.symform.com/>, Accessed: 2020.
- [15] Tor. <https://www.torproject.org/>, Accessed: 2020.
- [16] I. Abraham and D. Dolev. Byzantine agreement with optimal early stopping, optimal resilience and polynomial complexity. In *STOC*, 2015.
- [17] M. Andrychowicz and S. Dziembowski. Distributed cryptography based on the proofs of work. *IACR*, 2014.
- [18] B. Awerbuch and C. Scheideler. Robust random number generation for peer-to-peer systems. In *PODC*, 2006.
- [19] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker. Low-resource routing attacks against tor. In *WPES*, 2007.
- [20] J. Behl, T. Distler, and R. Kapitza. Hybrids on steroids: Sgx-based high performance bft. In *EuroSys*, 2017.
- [21] M. Ben-Or, B. Kelmer, and T. Rabin. Asynchronous secure computations with optimal resilience. In *PODC*, 1994.
- [22] P. Berman and J. A. Garay. Cloture votes: $n/4$ -resilient distributed consensus in $t+1$ rounds. *STOC*, 1993.
- [23] P. Berman and J. A. Garay. Randomized distributed agreement revisited. In *FTCS*, 1993.
- [24] P. Berman, J. A. Garay, and K. J. Perry. Optimal early stopping in distributed consensus. In *WDAG*, 1992.
- [25] N. Borisov, G. Danezis, P. Mittal, and P. Tabriz. Denial of service or denial of security? In *CCS*, 2007.
- [26] G. Bracha. An asynchronous $[(n-1)/3]$ -resilient consensus protocol. In *PODC*, 1984.
- [27] G. Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 1987.
- [28] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. *JACM*, 1985.
- [29] C. Cachin, K. Kursawe, F. Petzold, and V. Shoup. Secure and efficient asynchronous broadcast protocols. In *CRYPTO*, 2001.
- [30] C. Cachin, K. Kursawe, and V. Shoup. Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *J. Cryptology*, 2005.
- [31] C. Cachin and J. A. Poritz. Secure intrusion-tolerant replication on the internet. In *DSN*, 2002.
- [32] C. Cachin and S. Tessaro. Asynchronous verifiable information dispersal. In *SRDS*, 2005.
- [33] R. Canetti and T. Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *STOC*, 1993.
- [34] M. Castro, B. Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, 1999.
- [35] T. D. Chandra and S. Toueg. Time and message efficient reliable broadcasts. In *WDAG*, 1990.
- [36] B.-G. Chun, P. Maniatis, S. Shenker, and J. Kubiatowicz. Attested append-only memory: Making adversaries stick to their word. In *OSR*, 2007.
- [37] M. Correia, N. F. Neves, and P. Verissimo. How to tolerate half less one byzantine nodes in practical distributed systems. In *SRDS*, 2004.
- [38] M. Correia, P. Verissimo, and N. F. Neves. The design of a cots real-time distributed security kernel. In *EDCC*, 2002.
- [39] V. Costan and S. Devadas. Intel SGX explained. *IACR*, 2016.
- [40] G. Cybenko. Dynamic load balancing for distributed memory multiprocessors. *JPDC*, 1989.

- [41] D. Dolev and H. R. Strong. Authenticated algorithms for byzantine agreement. *SICOMP*, 1983.
- [42] P. Feldman and S. Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SICOMP*, 1997.
- [43] M. Fitzi and J. A. Garay. Efficient player-optimal protocols for strong and differential consensus. In *PODC*, 2003.
- [44] J. A. Garay and Y. Moses. Fully polynomial byzantine agreement in $t+1$ rounds. In *STOC*, 1993.
- [45] J. A. Garay and Y. Moses. Fully polynomial byzantine agreement for processors in rounds. *SICOMP*, 1998.
- [46] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *EUROCRYPT*, 1999.
- [47] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure distributed key generation for discrete-log based cryptosystems. *J. Cryptology*, 2007.
- [48] R. Guerraoui, F. Huc, and A.-M. Kermarrec. Highly dynamic distributed computing with byzantine failures. In *PODC*, 2013.
- [49] V. Hadzilacos and S. Toueg. Fault-tolerant broadcasts and related problems. In *Distributed Systems (2nd Ed.)*, 1993.
- [50] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg. Eclipse attacks on bitcoin’s peer-to-peer network. In *USENIX Security*, 2015.
- [51] Y. Jia, S. Tople, T. Moataz, D. Gong, P. Saxena, and Z. Liang. Robust p2p primitives using sgx enclaves. *IACR*, 2020.
- [52] B. M. Kapron, D. Kempe, V. King, J. Saia, and V. Sanwalani. Fast asynchronous byzantine agreement and leader election with full information. *TALG*, 2010.
- [53] J. Katz and C.-Y. Koo. On expected constant-round protocols for byzantine agreement. In *CRYPTO*, 2006.
- [54] J. Katz and Y. Lindell. *Introduction to modern cryptography*. 2014.
- [55] V. King and J. Saia. Scalable byzantine computation. *SIGACT News*, 2010.
- [56] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *TOPLAS*, 1982.
- [57] S. Lee, M.-W. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado. Inferring fine-grained control flow inside sgx enclaves with branch shadowing. *arXiv*, 2016.
- [58] D. Levin, J. R. Douceur, J. R. Lorch, and T. Moscibroda. Trinc: Small trusted hardware for large distributed systems. In *NSDI*, 2009.
- [59] Y. Lindell, A. Lysyanskaya, and T. Rabin. On the composition of authenticated byzantine agreement. *JACM*, 2006.
- [60] J. Liu, W. Li, G. O. Karame, and N. Asokan. Scalable byzantine consensus via hardware-assisted secret sharing. *arXiv*, 2016.
- [61] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena. A secure sharding protocol for open blockchains. In *CCS*, 2016.
- [62] S. Matetic, M. Ahmed, K. Kostianen, A. Dhar, D. Sommer, A. Gervais, A. Juels, and S. Capkun. Rote: Rollback protection for trusted execution. *IACR*, 2017.
- [63] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song. The honey badger of bft protocols. In *CCS*, 2016.
- [64] Ming-Wei-Shih, S. Lee, T. Kim, and M. Peinado. T-sgx: Eradicating controlled-channel attacks against enclave programs. 2017.
- [65] T. Moran and M. Naor. Split-ballot voting: everlasting privacy with distributed trust. *TISSEC*, 2010.
- [66] A. Mostefaoui, H. Moumen, and M. Raynal. Signature-free asynchronous Byzantine consensus with $t < \frac{n}{3}$ and $O(n^2)$ messages. In *PODC*, 2014.
- [67] K. Nayak, S. Kumar, A. Miller, and E. Shi. Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In *EuroS&P*, 2015.
- [68] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa. Oblivious multi-party machine learning on trusted processors. In *USENIX Security*, 2016.
- [69] P. R. Parvédy and M. Raynal. Optimal early stopping uniform consensus in synchronous systems with process omission failures. In *SPAA*, 2004.
- [70] R. Pass, E. Shi, and F. Tramèr. Formal abstractions for attested execution secure processors. *IACR*, 2016.
- [71] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 1980.
- [72] K. J. Perry and S. Toueg. Distributed agreement in the presence of processor and communication faults. *TSE*, 1986.
- [73] M. O. Rabin. Randomized byzantine generals. In *FOCS*, 1983.

- [74] M. O. Rabin. Transaction protection by beacons. *JCSS*, 1983.
- [75] A. Rao, K. Lakshminarayanan, S. Surana, R. Karp, and I. Stoica. Load balancing in structured p2p systems. In *IPTPS*, 2003.
- [76] S. Shinde, Z. L. Chua, V. Narayanan, and P. Saxena. Preventing page faults from telling your secrets. In *ASIACCS*, 2016.
- [77] E. Syta, P. Jovanovic, E. Kokoris-Kogias, N. Gailly, L. Gasser, I. Khoffi, M. J. Fischer, and B. Ford. Scalable bias-resistant distributed randomness. *IACR*, 2016.
- [78] V. Vaikuntanathan. *Randomized algorithms for reliable broadcast*. PhD thesis, 2009.
- [79] G. S. Veronese, M. Correia, A. N. Bessani, L. C. Lung, and P. Verissimo. Efficient byzantine fault-tolerance. *TC*, 2013.
- [80] Y. Xu, W. Cui, and M. Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *Security and Privacy*, 2015.

A Primitives and Formal Definitions

In this section, we first start by formally defining the syntax of the communication protocol between two peers, that we denote by Peer channel. Using this definition, we next define various failure modes and primitives. Using SGX, we assume that *execution integrity (P1)* is enforced. We then show that the following properties: *message integrity & authenticity (P2)*, *blind-box computation property (P3)* can be emulated based on the Blinded channel, executing on a particular program. Then we go ahead and formally define the *halt-on-divergence (P4)* property for any program running between two peers. Finally, we show how to reduce the byzantine model to a model where a peer can only replay, omit and delay, dubbed ROD, given that a Blinded channel exists.

Abstractly, a peer can be considered as the composition of two entities: an Enclave and an OS. The OS models the untrusted entity including the operating system and memory. It has access to all the system resources such as file system, network and others. The OS can arbitrarily invoke an enclave program and start its execution. The Enclave models the isolated memory space that loads the program and executes it securely. Thus, Enclave corresponds to the trusted entity of a peer. A concurrent work provides a formal study to show that SGX enclaves can be considered as a trusted entity [70]. The Enclave of the two Peers can interact with each other via their OSs. We formally define a Peer channel as a protocol, $Peer^{ch}$, between a sender $Peer_s = (Enclave_s, OS_s)$ and a receiver $Peer_r = (Enclave_r, OS_r)$. A Peer channel can be seen as a generalization of the traditional secure communication

channel between two parties. The main difference is that the definition of $Peer^{ch}$ protocol is augmented with the program π running within the trusted Enclave.

We define four progressively stronger failure modes: *honest*, *general omission*, ROD and *byzantine* modes of $Peer^{ch}$. Here we introduce a ROD model as an intermediate model, wherein the adversary can only a) Replay b) Omit c) or Delay messages during a protocol, or follow it as prescribed.

We define two primitives: a) Blinded channels and b) halt-on-divergence. Informally, a Blinded channel guarantees confidentiality and integrity of a message over a Peer channel $Peer^{ch} = (Init, Write, Transfer, Read)$.

We show how we build a $Peer^{ch}$ channel using SGX where $Enclave_s$ and $Enclave_r$ are trusted entities. A $Peer^{ch}_{sgx}$ channel is a Blinded channel, and therefore enforces both (P2) and (P3) properties. In particular, we consider that there is a $KeyEx_{\pi}$ protocol between $Enclave_s$ and $Enclave_r$ that is used to generate a session key for a program π . Whenever there is a new program the key has to be re-generated. The key exchange protocol can be instantiated using Diffie-Hellman key exchange, referring to [54] Chapter 9. We use SGX remote attestation to verify that both parties run their code inside an Enclave. While this step is neither required nor captured in the $Peer^{ch}$ definition, it is mandatory to guarantee our *execution integrity (P1)*. Due to space constraints, we defer the formalization and proof to Appendix A in the full version [51].

B ERB Analysis

In this section, we use the same terminology used in Appendix A, namely, we assume that between any two nodes of the network, an $Peer^{ch}_{sgx}$ instantiation of the Blinded Peer channel is enabled. In particular, it provides us with both *message Integrity & authenticity (P2)* and *blind-box computation (P3)* properties. Throughout this section, we implicitly consider that the program is publicly available, and therefore its *execution integrity (P1)* is enforced.

Theorem B.1. *If $N \geq 2t + 1$ where t is the upper bound on the number of byzantine peers, and $Peer^{ch}_{sgx}$ is a Blinded Peer channel, then ERB is a reliable broadcast protocol as defined in Definition 2.1 with worst-case round complexity equal to $t + 2$ and communication complexity equal to $O(N^2)$.*

Proof. Due to space limitation, we defer the proof of the five requirements of terminating reliable broadcast to Appendix C in the full version [51]. Note that the assumption that the peers communicate using $Peer^{ch}_{sgx}$ implies that a byzantine node can only delay, omit or replay messages. As long as the network is synchronous with a fixed time interval for a round to complete, delaying is then equivalent to omitting a message, as the message will not be considered by honest nodes past the round, enforcing therefore the *lockstep execution (P5)* property. Replaying a message is also ineffective as every peer is identified by a sequence number as well,

that is generated by the trusted enclave in the Peer channel, and therefore enforcing the *message freshness* ($P6$) property. Under the assumption that $\text{Peer}_{\text{sgx}}^{\text{Ch}}$ is a Blinded channel, we can replace all occurrences of Multicast by communication between two trusted parties. To sum up, and throughout the proof, it is valid to consider that if there is a delay, omission or replay, this will be equivalent to considering that the first party does not send any message. \square

C Unoptimized ERNG Analysis

In this section, similar to Appendix A, we denote by the ROD mode, a mode where peers in a network \mathcal{P} can only replay, omit and delay messages.

Theorem C.1. *If \mathcal{P} operates in the ROD mode, then the bias of G $\beta(G) = 1$.*

Proof. Note that while G can be modeled as a multi-variate function, it does not capture the sequencing of inputs. For our proof to go through, we need to first show that the sequencing of ERNG is guaranteed and a node can only participate with its input if it starts synchronously with all nodes. For this, we have the following two cases:

- *early start:* if a byzantine node transmits its INIT at $\text{rnd} = 1$, the node outputs (either m or \perp) will be considered as an input for G ,
- *late start:* if a byzantine holds the INIT message until seeing the output, then its input will not be added to S_{final} as the message will be considered delayed. The output of G in this case equals \perp

Note that for both cases, the nodes have to start the protocol at $\text{rnd} = 1$ if they want to participate with their inputs in the final output. Moreover, based on the Blinded channel, we know that nodes can only obtain the final output of G while not viewing any internal state of G , which enforce the *blind-box computation* ($P3$) property. That is, it is valid to consider G as a multi-variate function that is fed all inputs at once. Let us denote by X the random variable that captures the output of G such that $X = X_1 \oplus \dots \oplus X_N$, where X_i 's are random variables that capture the input provided by every node in \mathcal{P} , for all $i \in [N]$. As \mathcal{P} operates in the ROD mode, all honest nodes receive the same set S_{final} at the end of the protocol. We then can rewrite X such that $X = \bigoplus_{i=1}^{\kappa} X_i \oplus \bigoplus_{i=\kappa+1}^N X_i$, where $\kappa = |S_{\text{final}}|$. In the following, we need to show that $E_G[S] = E[S] = \frac{|S|}{2^k}$, for all $S \subseteq \{0, 1\}^k$. Note that $E_G[S] = \Pr[X \in S]$, and therefore it is sufficient to compute $\Pr[X \in S]$.

The second equality follows from the fact that all events are disjoint. Now for a given $x \in S$, $\Pr[X = x] = \Pr[\bigoplus_{i=1}^{\kappa} X_i \oplus \bigoplus_{i=\kappa+1}^N X_i = x] = \frac{1}{2^{s \cdot k}} |\{x_2, \dots, x_{\kappa} \in \{0, 1\}^k\}| = \frac{1}{2^s}$. Thus, $\Pr[X \in S] = \frac{|S|}{2^s}$. This concludes our proof. \square

Algorithm 4: ERNG: Enclaved unbiased random number generation protocol executed by peer p_i .

Input: A P2P network \mathcal{P} composed of N nodes

Output: A unbiased random number r

```

• initialization:  $S_M \leftarrow \emptyset; S_{\text{final}} \leftarrow \emptyset; S_{\text{chosen}} \leftarrow \emptyset; \text{rnd} \leftarrow 1$ 
for  $\text{rnd} \leq \gamma + 4$  do
  if  $\text{rnd} = 1$  then
    • every peer  $p_i$  compute  $r_i \xleftarrow{\$} \{0, \dots, \frac{N}{2\gamma} - 1\}$ ;
    if  $r_i = 0$  then
      Multicast( $\text{id}_i, \text{val}$ ), where
       $\text{val} = \langle \text{CHOSEN}, \text{id}_i, \text{seq}_j, \perp, 1 \rangle$ ;
       $S_{\text{chosen}} \leftarrow \{\text{id}_i\}$ ;
    end
    • upon receiving  $\text{val} = \langle \text{CHOSEN}, \text{id}_j, \text{seq}_j, m_j, \text{rnd}_j \rangle$ 
      if type = CHOSEN and  $\text{rnd}_j = 1$  and  $\text{seq}_j = \text{seq}_i$  then
         $S_{\text{chosen}} \leftarrow S_{\text{chosen}} \cup \{\text{id}_j\}$ ;
      end
    end
  • if  $r_i = 0$  and  $\text{rnd} = 2$  then
    compute  $r'_i \xleftarrow{\$} \{0, \dots, \gamma - 1\}$ ;
    if  $r'_i = 0$  then
      initiate ERB with inputs  $m_i \xleftarrow{\$} \{0, 1\}^k$ ,  $\text{seq}_i$  and
      peers in  $S_{\text{chosen}}$ ;
    end
     $\text{seq}'_j \leftarrow \text{seq}_j$ , for all  $\text{id}_j \in S_{\text{chosen}}$ ;
    end
  if  $r_i = 0$  and  $3 \leq \text{rnd} \leq \gamma + 2$  then
    execute ERB instances and wait for the output;
  end
  if  $r_i = 0$  and  $\text{rnd} = \gamma + 3$  then
    Wait( $\text{rnd}$ ) then obtain  $M_i = \{\hat{m}_1, \dots, \hat{m}_i\}$ ;
     $\text{seq}'_j \leftarrow \text{seq}'_j$ , for all  $\text{id}_j \in S_{\text{chosen}}$ ;
  end
  if  $\text{rnd} = \gamma + 4$  then
    • if  $r_i = 0$  then
       $S_M \leftarrow S_M \cup \{M_i\}$ ;
      Multicast( $\text{id}_i, \langle \text{FINAL}, \text{id}_i, M_i, \text{seq}_j, \gamma + 4 \rangle$ );
    end
    • upon receiving  $\text{val} = \langle \text{FINAL}, M_j, \text{seq}'_j, \text{rnd}_j \rangle$ 
      if  $\text{rnd}_j = \gamma + 4$  and  $\text{seq}'_j = \text{seq}_j$  then
         $S_M \leftarrow S_M \cup \{M_j\}$ ;
        if # of  $M_{\kappa} \geq \gamma + 1$  where  $M_{\kappa} \in S_M$  then
           $S_{\text{final}} \leftarrow M_{\kappa}$ ;
          accept  $r = \bigoplus_{v \in S_{\text{final}}} v$ .
        end
      end
    end
  end
   $\text{rnd} \leftarrow \text{rnd} + 1$ ;
end
•  $\text{seq}_j \leftarrow \text{seq}_j + 1$ , for all  $j \in [N]$ ;

```

D Optimized ERNG

We present a pseudo-solution of our optimized ERNG in Algorithm 4.

Lemma D.1. *If up to $t = \frac{N}{3}$ nodes are byzantine, then with at least $1 - \text{negl}(\gamma)$ probability, the representative cluster has more than γ honest nodes, and less than γ byzantine nodes.*

Proof. In ERNG at round 1, every node picks uniformly at random a value from $\{0, \dots, \frac{N}{2\gamma} - 1\}$. That is, every node has a probability equal to $q = \frac{2\gamma}{N}$ to be chosen as a representative.

Let H_i and B_i be two random variable that equal 1 if the i^{th} honest and byzantine node is chosen respectively, otherwise they equal zero. Let us denote by $H = \sum_{i=1}^{2t} H_i$ and $B = \sum_{i=1}^t B_i$ the number of selected honest and byzantine nodes in the cluster. Then both H and B are distributed following a binomial distribution with a number of trials equal to $2t$ and t , respectively. We have $E[H] = \sum_{i=1}^{2t} E[H_i] = 2t \cdot \frac{2\gamma}{N} = \frac{4t\gamma}{N}$. Similarly, $E[B] = \frac{2t\gamma}{N}$. Based on two variations of Chernoff bound, considering $t = \frac{N}{3}$, we obtain that

$$\Pr[H > (1 - \delta_1) \frac{4\gamma}{3}] \geq 1 - e^{-\frac{2\delta_1^2\gamma}{3}},$$

similarly, $\Pr[B < (1 + \delta_2) \frac{2\gamma}{3}] \geq 1 - e^{-\frac{2\delta_2^2\gamma}{9}}$, where $\delta_1, \delta_2 < 1$. For a choice of $\delta_1 = \frac{1}{4}$ and $\delta_2 = \frac{1}{3}$, we obtain,

$$\Pr[H > \gamma] \geq 1 - e^{-\frac{\gamma}{24}},$$

and,

$$\Pr[B < \gamma] \geq 1 - e^{-\frac{\gamma}{4t}}.$$

□

Lemma D.2. *If $\gamma' = \sqrt{\gamma}$, then the probability that $\Omega(\sqrt{\gamma})$ honest nodes are selected to be in the second representative cluster is at least $1 - \text{negl}(\gamma)$.*

Proof. Based on Algorithm 4, every node in the cluster has a probability of $\frac{1}{\gamma}$ to be chosen. Let us denote by X_i the random variable equal to one if the node is selected. We then denote by, $H' = \sum_{i=1}^H X_i$, the random variable that counts the number of honest node in the second cluster. Based on Wald's equation, we obtain $E[H'] = \frac{E[H]}{\gamma} = \frac{4\gamma}{3\gamma}$. Then, based on Chernoff bound, we obtain for $\delta < 1$,

$$\Pr[H' > (1 - \delta) \cdot \frac{4\gamma}{3\gamma}] \geq 1 - e^{-\frac{4\delta^2\gamma}{3\gamma}}$$

if we set $\delta = 1 - \frac{1}{\gamma}$ and $\gamma' = \sqrt{\gamma}$, then we obtain

$$\Pr[H' > \frac{4\sqrt{\gamma}}{3}] \geq 1 - e^{-\sqrt{\gamma}}.$$

This ends out proof. □

Note that we can obtain better bounds if we consider computing the pmf of H' as it follows a binomial distribution with a binomial number of trials

Corollary D.1. *If $\gamma' = \sqrt{\gamma}$, then the size of the first and second representative clusters is in $O(\gamma)$ and $O(\sqrt{\gamma})$ w.h.p*

The proof of the corollary directly follows from Lemma D.2.

Theorem D.1. Agreement: *All honest nodes eventually agree on the same common set S_{final} in ERNG.*

Proof. In round 1, $|S_{\text{chosen}}|$ nodes are uniformly at random selected to be part of the representative cluster. Based on Lemma D.1, we have shown that the cluster contains *strictly* more than γ honest nodes, and *strictly* less than γ byzantine nodes w.h.p. when $t < \frac{N}{3}$. That is, we have created a new smaller P2P network S_{chosen} in which the honest nodes represent the majority. In the cluster, all honest nodes know each other, but byzantine nodes may deliberately not contact honest nodes on purpose. In this case, the cluster will be more robust with less byzantine nodes. Thus, all the results introduced for ERB will hold for this cluster of nodes.

From round 2 to round $\gamma + 3$, the second cluster has more than $\sqrt{\gamma}$ honest nodes w.h.p. according to Lemma D.2. For each instance of ERB— whether initiated by an honest or byzantine node, the honest representative nodes will agree on a same message. Since there is at least one honest sender, all honest nodes will accept the honest sender's message for its run of ERB. After around $O(\sqrt{\gamma})$ runs, all honest nodes will agree on the same set of random numbers. Since the number of honest representative nodes is larger than γ and all of them will multicast FINAL messages for the same set of messages in round $\gamma + 4$, then all honest nodes will receive adequate FINAL messages to accept the common set S_{final} . □

Theorem D.2. Unbiasedness: *The output of the ERNG protocol is an unbiased random number.*

sketch. Given Theorem D.1, we know that all honest nodes agree on the same set S_M . On the other hand, leveraging Peer_{sgx}^{Ch} Peer channel, we know that all random numbers in the ERNG protocol are generated within the SGX enclave and never tempered with as the network is in the ROD model. Finally, it is sufficient to show that if all random numbers generated in SGX are random then the output of ERNG is an unbiased random number, which holds given SGX primitive generates unbiased random number against the operating system according to Theorem C.1. □

In ERNG, since the message m_i is a random number generated by the SGX and proposed by the peer p_i , then eventually every honest node accepts the same set S_{final} of random numbers according to Theorem D.1. By performing exclusive disjunction (or XOR) of all the random numbers in S_{final} , every honest node can obtain a common random number r . In the meantime, the random number r is unbiased against byzantine nodes.