



EESMR: Energy Efficient BFT—SMR for the masses

Adithya Bhat
abhatk@purdue.edu
Purdue University
West Lafayette, IN, USA

Akhil Bandrupalli
abandaru@purdue.edu
Purdue University
West Lafayette, IN, USA

Manish Nagaraj
mnagara@purdue.edu
Purdue University
West Lafayette, IN, USA

Saurabh Bagchi
sbagchi@purdue.edu
Purdue University
West Lafayette, IN, USA

Aniket Kate
aniket@purdue.edu
Purdue University
West Lafayette, IN, USA

Michael K. Reiter*
michael.reiter@duke.edu
Duke University
Durham, NC, USA

ABSTRACT

Modern Byzantine Fault-Tolerant State Machine Replication (BFT-SMR) solutions focus on reducing communication complexity, improving throughput, or lowering latency. This work explores the energy efficiency of BFT-SMR protocols. First, we propose a novel SMR protocol that optimizes for the steady state, i.e., when the leader is correct. This is done by reducing the number of required signatures per consensus unit and the communication complexity by order of the number of nodes n compared to the state-of-the-art BFT-SMR solutions. Concretely, we employ the idea that a quorum (collection) of signatures on a proposed value is avoidable during the failure-free runs. Second, we model and analyze the energy efficiency of protocols and argue why the steady-state needs to be optimized. Third, we present an application in the cyber-physical system (CPS) setting, where we consider a partially connected system by optionally leveraging wireless multicasts among neighbors. We analytically determine the parameter ranges for when our proposed protocol offers better energy efficiency than communicating with a baseline protocol utilizing an external trusted node. We present a hypergraph-based network model and generalize previous fault tolerance results to the model. Finally, we demonstrate our approach's practicality by analyzing our protocol's energy efficiency through experiments on a CPS test bed. In particular, we observe as high as 64% energy savings when compared to the state-of-the-art SMR solution for $n = 10$ settings using BLE.

CCS CONCEPTS

• Security and privacy → Distributed systems security.

KEYWORDS

Energy Efficient, State Machine Replication, Byzantine Fault Tolerant

ACM Reference Format:

Adithya Bhat, Akhil Bandrupalli, Manish Nagaraj, Saurabh Bagchi, Aniket Kate, and Michael K. Reiter*. 2023. EESMR: Energy Efficient BFT—SMR for the masses. In *24th International Middleware Conference (Middleware '23)*, December 11–15, 2023, Bologna, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3590140.3592848>

1 INTRODUCTION

Byzantine faults are not new to the computing world. For more than four decades, the distributed system literature has addressed Byzantine faults using a variety of consensus mechanisms over replicated nodes. Among those mechanisms, we focus on State Machine Replication (SMR) [6, 15, 21, 39, 48, 61, 63, 76] as it is a generic consensus problem that ensures that the correct nodes agree on the same sequence of values in the presence of Byzantine nodes. SMR is of general interest in applications such as Multi-party computation [9, 13] which require repeated usages of one-shot consensus problems such as Byzantine Agreement (BA) [49] or Reliable Broadcast [19]. An SMR can also implement a distributed ledger (or a blockchain), which has resulted in the tremendous recent interest in these protocols [4, 6, 21, 24, 39, 76].

Due to its widespread adoption, energy-efficiency of SMR is an important factor as they are employed in data-centers to ensure system resilience [1, 65]. They are also considered for the CPS/IoT environments [17, 26, 50, 56–58, 60, 77–79] where energy-efficiency can determine how long the system functions. Even in energy-rich environments such as data-centers, energy costs *can be reduced* by running a more energy-efficient SMR protocol in large scales, and thus this problem is of general interest.

Meanwhile, Cyber-Physical Systems (CPS) provide integrated computational, mechanical, and communication components that interact with each other. CPS can vary from a system of temperature sensors deployed in power plants [62] to a network of soil nutrient monitoring sensors spread over a farm field, to unmanned aerial vehicles or drones deployed to monitor traffic [67]. In many of these scenarios, the various nodes or devices in such systems need to agree on a common state. For example, in a military setting, several CPS nodes could be deployed to survey an area; they should maintain a shared state, which they report during their sporadic contacts with a base station.

*The author was also affiliated with Chainlink Labs at the time of this work.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
Middleware '23, December 11–15, 2023, Bologna, Italy
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0177-1/23/12.
<https://doi.org/10.1145/3590140.3592848>

An adversary may try to intervene in these consensus processes by compromising some nodes: the compromised nodes may communicate incorrect values, possibly in collusion with other compromised nodes, or prevent the correct nodes in the system from performing a consistent action (e.g., Byzantine Generals Problem [49]). While such adversarial threats are obvious in battlefield CPS, they are also becoming pervasive across the entire CPS application spectrum [17, 26, 50, 77–79]. For example, a recent US Department of Homeland Security (DHS) report on precision agriculture [54] asks to pay particular attention to Byzantine/malicious faults. The agency observes that an intentional falsification of data can disrupt crop or livestock sectors as the introduction of rogue data into a sensor network can damage crops or herds.

Today, with the focus on Internet-based instantiations, the state-of-the-art SMR solutions try to solely reduce the latency or communication overhead by using, for example, threshold signature schemes [18]. While theoretically interesting, threshold signatures are computationally intensive and have high energy demands, making SMR unsuitable for CPS.¹ In general, naïvely employing an SMR protocol that is not optimized for energy efficiency reduces the lifetime of CPS, and we set out to explore energy efficiency for the SMR problem for the CPS setting.

Indeed, the state-of-the-art latency-optimal SMR protocols OptSync [63] and Sync HotStuff [6] are not energy-efficient solutions for SMR due to their high (bit-level) communication complexity and extensive use of expensive cryptographic primitives. In general, for SMR in CPS, since communication and cryptography are the main sources of energy drain, one might try to use the state-of-the-art SMR protocol, and optimize away the bottleneck. However, generally optimizing in one direction leads to degradation of the other. For example, an attempt toward energy efficiency can be to substitute expensive cryptography such as threshold signatures or public-key cryptography with simpler symmetric-key primitives. Doing so results in increased communication footprint [11, 14, 21] or reduced fault tolerance [10, 21, 52], both of which are not ideal. **Contributions.** This work makes the following four contributions towards energy-efficient consensus for CPS:

(1) We present EESMR², an energy-efficient, leader-based SMR protocol (Section 3). In the steady state, i.e., when the leader is correct, EESMR consumes less energy than the state-of-the-art SMR protocols Sync HotStuff and OptSync. The trade-off is that during the view-change phase, i.e., the sub-protocol to change the leader, EESMR performs slightly worse than Sync HotStuff [6] by adding an extra round. Concretely, EESMR avoids using computationally expensive certificates and voting in the steady state and pushes those overheads to the view-change phase invoked when a leader stops making progress or if it equivocates.

(2) As every application setting is different, we derive an easy-to-use template for comparing SMR protocols in various CPS settings (Section 4). We model the energy cost of protocols as functions of system parameters such as the number of nodes n and message size m . Our analysis shows the conditions (number of faults, number

of nodes, wireless communication modalities, etc.) where an SMR protocol A will be more energy-efficient than a competing protocol B . We demonstrate the use of this analysis to determine the conditions EESMR is desirable compared to a baseline solution where nodes always communicate with a trusted control node.

(3) We observe that in some CPS settings, a node can communicate with its k neighbors using a single multicast message, instead of k unicasts with each neighbor. We take advantage of such multicasts if they are available, reliable and are more energy-efficient when handling protocol messages than unicasts, to improve the energy efficiency of protocols. In order to capture this communication modality, we extend the standard directed graph model of the network using hypergraphs.

(4) Finally, using an experimental CPS test-bed, we analyze the energy efficiency of cryptographic primitives along with parameters as well as the available communication modalities (Section 5). Moreover, we measure and choose energy-efficient digital signature schemes. As EESMR and other SMR protocols heavily employ the communication pattern of a single leader signing and the others verifying, we find that the verification-efficient RSA signatures are more energy-efficient than the ECDSA signature scheme. We also analyze the energy-vs.-reliability trade-off for Bluetooth Low Energy (BLE) multicasts and demonstrate that k -casts with four nines reliability are more efficient than unicasts. For our optimal choice of cryptographic primitives and communication modality, we show that EESMR is 2.8 times more energy-efficient than the state-of-the-art SMR protocol Sync HotStuff [6] in the failure-free runs, and 2 times worse than Sync HotStuff during leader changes, demonstrating the positive trade-offs of our approach.

2 PRELIMINARIES

System Model. Consider an n -node system $\mathcal{N} = \{p_1, \dots, p_n\}$, with up to $f < n/2$ nodes being Byzantine. We assume a static fully-connected point-to-point communication graph. The links are bounded synchronous [4, 6, 8, 35, 38, 59, 73, 74, 80]. We assume that Δ is a public upper bound on the message delivery time for any correct sender. The adversary can control the delivery of messages as long as the bounded-synchrony assumption is not violated.

State Machine Replication — SMR. SMR is an abstraction of a state machine that applies input requests from clients to a state and outputs new state. An SMR protocol [61] (defined formally in Definition 2.1) ensures all correct nodes implement the SMR abstraction where the nodes reach agreement on an ordering of requests (or transactions) via a linearizable log.

Definition 2.1 (State Machine Replication (SMR) [6]). An SMR protocol generates a linearizable log of transactions satisfying the following:

- (1) **Safety.** For any position m in the log, if two correct nodes output B_m and B_m^* respectively, then $B_m = B_m^*$.
- (2) **Liveness.** Each client request is eventually committed by all the correct nodes.

We make the same assumptions as existing SMR literature [6, 23, 24, 57, 58, 63] with respect to the role of clients and the validity of the requests. Some example assumptions include: all the clients are honest when measuring the throughput, every transaction can be validated for correctness, every protocol message has an identifier

¹Threshold signing operations are also inefficient that many of the protocols, e.g., [6, 39, 76], forgo threshold signatures in their implementations, opting instead for multiple traditional digital signatures and resulting linear growth in communication complexity.

²EESMR stands for *Energy Efficient State Machine Replication*. Also pronounced as easy-SMR.

for that instance so messages from one instance will not become equivocation³ in another instance, the hash functions are all seeded, PKI is used to setup (possibly threshold) keys before starting the protocol, and the public information is agreed upon by all the nodes as part of the setup before the start of the protocol.

Blocks. In our instantiation of SMR, we use blocks to denote a unit of the linearizable log. A block is defined as

$block.contents = \text{CMDS},$
 $block.parent = \text{hash of the parent block}.$

$block.parent$ is the hash of the *parent* block. $block.contents$ is the set of *requests* from clients called CMDS. We let G be the genesis block with height 0, and recursively define the height of other blocks as the height of the parent plus one. We refer to the *highest block* as the block with the highest height.

Clocks. We assume that the clocks of all the nodes have a constant bounded skew achieved using network time protocol [64, 66, 68] and that all the nodes start within time Δ of each other, which is ensured using a clock synchronization protocol [4]. All nodes start at local time $t = 0^4$.

Cryptographic Primitives. Generally, there are two classes of authentication primitives: digital signatures which are asymmetric (public-key) primitives, and Message Authentication Codes (MACs) schemes which are symmetric-key primitives.

A digital signature scheme consists of three algorithms which fail with negligible probability in the security parameter: KeyGen, Sign, and Verify. $(sk, pk) \leftarrow \text{KeyGen}(\kappa)$, where κ is the security parameter. sk is the secret key used to sign the message and pk is the public key used by others to verify that the message was signed by whoever possesses sk . $\sigma \leftarrow \text{Sign}(sk, m)$, where σ denotes the signature on m signed by using secret key sk . $\{0, 1\} \leftarrow \text{Verify}(m, \sigma, pk)$ outputs 1 if the signature σ signed by sk on m is valid, and outputs 0 otherwise. We denote $\langle m \rangle_i$ as (m, σ) where $\sigma \leftarrow \text{Sign}(sk_i, m)$ and sk_i is the secret key of node p_i .

A Message Authentication Code (MAC) scheme consists of three similar algorithms. A MAC is generally more energy-efficient than a digital signature, but a digital signature scheme ensures transferable authentication, i.e., a node can convince others that a particular node sent a particular message using a signature on it without possessing sk which is useful for detecting equivocation. In a MAC scheme, it is harder to verify equivocation [11] in the distributed environment even for a moderate number of nodes.

We use both primitives for energy-analysis, and use digital signatures in our protocol. We assume a public key infrastructure (PKI) to set up digital signatures. Similar to the related work [4, 6, 63], threshold signatures can reduce the certificate size to $O(1)$.

3 EESMR PROTOCOL

The EESMR protocol implements the SMR abstraction (Definition 2.1). It builds a chain of blocks linked using hashes, and all the nodes execute the client requests in the blocks in the same sequence. The clients wait to receive $f + 1$ identical acknowledgments with

³An equivocation is a Byzantine behavior where a node sends two conflicting messages to two different nodes.

⁴The local time may be off by Δ across different nodes. This is known as the non-lock step model [6] as opposed to the lock-step synchrony model which assumes that all the nodes are in the same round at the same time.

execution results and accepts the results. We omit the clients and the execution from the rest of the discussion and focus on the nodes that run the SMR.

EESMR protocol runs in a sequence of *rounds* and *views* with monotonically increasing numbers from 1. Each view consists of a dedicated unique leader $L \in \mathcal{N}$ known to all. Within a view, the leader proposes blocks in every round. All nodes maintain pending commands in a local data structure $txpool$. The leader proposes blocks using the commands from $txpool$ and the other nodes on committing a block, remove the commands in the block from the $txpool$.

3.1 Data Structures

Messages and Quorum Certificates. Every message m contains the view number in the field $m.view$ and a threshold signature validating it in the field $m.dataSig$ and $m.viewSig$. The message is customizable into different messages in the protocol by specifying its type in $m.type$ (as specified in Algorithm 1). $n/2 + 1$ valid threshold signed messages m from the same view and type is combined into a *Quorum Certificate* (using the QC function in Algorithm 1).

Algorithm 1 Helper functions for EESMR (for node p_i).

```

101: function MSG( $type, data$ )
102:    $m.type \leftarrow type$ 
103:    $m.data \leftarrow data$ 
104:    $m.view \leftarrow v_{cur}$ 
105:    $m.viewSig \leftarrow \langle m.type, v_{cur} \rangle_i$ 
106:    $m.dataSig \leftarrow \langle m.data, v_{cur} \rangle_i$ 
107:   return  $m$ 
108: function CREATEPROPOSAL( $block, CMDS$ )
109:    $newBlock.parent \leftarrow block$ 
110:    $newBlock.contents \leftarrow CMDS$ 
111:   return  $newBlock$ 
112: function MATCHINGMSG( $m, type, view$ )
113:   return  $m.type = type \wedge m.view = view$ 
114: function QC( $V$ )
115:    $qc.type \leftarrow m.type : m \in V$ 
116:    $qc.view \leftarrow m.view : m \in V$ 
117:    $qc.cert \leftarrow C(m.type, m.view) : \{m.viewSig \mid m \in V\}$ 
118:   return  $qc$ 
119: function MATCHINGQC( $qc, type, view$ )
120:   return  $qc.type = type \wedge qc.view = view$ 
121: function LOCKCOMPARE( $B_{lck}, b$ )
122:   if  $view(B_{lck}) \neq view(b)$  and  $b$  extends  $B_{lck}$  then
123:     return  $b$ 
124:   if  $round(B_{lck}) \neq round(b)$  and  $b$  extends  $B_{lck}$  then
125:     return  $b$ 
126:   return  $B_{lck}$  // Retain the old lock

```

Book-keeping variables. We maintain state variables to track the current state of the protocol: (i) v_{cur} tracks the current view number, initialized with 1. (ii) r_{cur} tracks the current round number, initialized with 3. (iii) L tracks the current leader. (iv) B_{lck} tracks the currently locked chain, initialized to the genesis block G . A locked chain is only updated using the LOCKCOMPARE function in Line 121 and *always* extends the highest committed block. A node will never allow any other chain to fork-away from the locked block B_{lck} (unless it is safe to do so, described in the view-change)

while ensuring that it always extends the highest committed block. (v) B_{com} tracks the highest committed block, and is initialized to the genesis block G . (vi) Timers $T_{commit}(block)$ and $T_{blame}(v_{cur})$ are used to commit blocks and send blame messages for a round respectively.

3.2 Protocol Overview

Our protocol is motivated from Sync HotStuff [6] and OptSync [63]. We present the technical version of our protocol in Algorithm 2. A crash-version of our protocol can be obtained by removing the blocks in Lines 220 and 224 as they are associated with equivocation which is allowed only for Byzantine faults.

The description contains code regions, which are atomically executed. All blocks such as the conditional **on** regions are running concurrently. The same node may act in multiple roles, for instance, a node maybe a leader as well as a node and thus execute both parts of the **as** regions.

The protocol consists of two sub-protocols: *steady-state protocol* and *view-change protocol*. The steady-state is a period when the leader is behaving correctly. The steady-state protocol can be a separate protocol, with safety guarantees but no liveness guarantees. The view-change is the phase that handles Byzantine behavior of the steady-state leader, by transitioning from view v with the Byzantine leader to view $v+1$. If the leader of view $v+1$ is Byzantine, then another view-change protocol is initiated. The view-change protocol needs to ensure both safety and liveness. Technically, with minor modifications, the view-change protocol can be a stand-alone SMR protocol. However, it is inefficient, and is only used to bring the system back into a steady-state to enjoy efficiency benefits of the steady-state.

The protocol proceeds in views $1, \dots$ and rounds $3, \dots$. The first 2 rounds in the counter are reserved for the view-change protocol, and ensures that the view starts with a safe block. We assume that the leaders are chosen using the $LEADER(v_{cur})$ function. This function can be round-robin for simplicity, but for expected constant-latency it is required that the leaders are chosen randomly [4]. All messages are validated for correctness, such as formatting, signature checks, etc, and processed when the round is current. If messages are received from rounds larger than r_{cur} they are buffered and processed on entering the round to ensure liveness.

Note on chain synchronization. Our description of the protocol accounts for chain-synchronization, i.e., when a node obtains a block and does not know its parent blocks, it will request them from the sender first. To ease the exposition, we do not include this request-response in the protocol description, but the timers in our description account for these synchronization. Note that the chain synchronization is not a by-product of our protocol, but is used in protocol and thus does not affect the results. Since a Byzantine node can trigger chain synchronization once every round, this results in a communication-complexity $O(n^2)$ per round.

3.3 EESMR — Steady State

In the steady-state, the leader proposes blocks for every round. These blocks are broadcast to all the nodes (Line 205). The leader continuously streams proposals allowing the protocol to use all available bandwidth. The other nodes receive the proposal before

the timer (Line 216) runs out. On receiving a new block, all nodes also update their locked block B_{lck} . All nodes broadcast this proposal to all the other nodes (Line 213).

Committing. For a block B to be committed Lines 214 and 278, the node waits for 4Δ time using $T_{commit}(B)$ after broadcasting B to all the nodes. During this wait, the node ensures that it does not hear any equivocating blocks for that view and round. When the $T_{commit}(B)$ expires, the node also tracks the highest committed block in the variable B_{com} .

The 4Δ wait ensures that if a node commits a block at time t , then the nodes must have forwarded it to all the nodes at time $t - 4\Delta$ which was received by all the correct nodes by time $t - 3\Delta$. After chain-synchronization, all nodes would have this block by time $t - \Delta$. Since, no equivocation was heard by time t , all the correct nodes will *lock* on this block, i.e., the B_{lck} of all correct nodes will always extend B . The view-change ensures that the first block in the new view always extends the highest committed block of all correct nodes, thus ensuring safety for B even if a view-change occurs. We discuss the details of this in the description of the view-change protocol.

Voting in the head. In the steady-state, by waiting for 4Δ time before committing a block B and not hearing any equivocating blocks, a node has implicitly collected votes for B from all the correct nodes. We describe this as *voting in the head*, as opposed to explicit quorum certificate construction. These implicit votes are made explicit during a view-change. On a high-level, in the view-change protocol, the correct nodes send their vote for blocks that extend their locked block B_{lck} . Due to hash-chaining of blocks, a vote for any child block is also a vote on all its ancestors, and thus all the committed blocks obtain explicit votes during the view-change.

Firstly, this results in $O(1)$ signing operations per node for every committed block in the steady-state, which contributes to energy-savings. Secondly, since this is a re-broadcast of the leader's proposal and not broadcast of explicit vote messages, in partially connected networks, it results in a flooding of only one message reducing the communication complexity of our protocol.

3.4 EESMR — View Change

In EESMR, a view change can occur if the leader of a view (i) allows T_{blame} to time out for at least 1 correct node, (ii) equivocates in a round with two blocks at the same height, or (iii) does not correctly extend the highest committed block during the view change (rounds 1 and 2).

If a correct node p_i does not receive the next proposal on time, it sends a blame message (Line 216). If no correct node hears a block then there will be $n - f > f$ blames, and a view change will occur. But, if a correct node p_i blames, it may not always result in a view-change, as less than $f + 1$ correct nodes may send this, but it is still safe. If other correct nodes hear a proposal for that round, then they will forward it to node p_i within Δ and this results in a reset of the blame timer (Line 210).

A leader L can equivocate by sending B and B' for any round (not just the latest round). From bounded-synchrony, we can ensure all correct nodes will hear it in time Δ . On observing two equivocating (conflicting) proposals (Line 220), all correct nodes $n - f > f$ send

Algorithm 2 EESMR Protocol (for node p_i).

```

201: for  $v_{cur} \leftarrow 1, 2, 3, \dots$  do
202:   for  $r_{cur} \leftarrow 3, 4, 5, \dots$  do
203:     ▶ Steady-State Phase (Safe)
204:     as a leader //  $L \leftarrow \text{LEADER}(v_{cur})$ 
205:       // CMDS is obtained from  $txpool$ 
206:        $newBlock \leftarrow \text{CREATEPROPOSAL}(B_{lck}, \text{CMDS})$ 
207:        $propData \leftarrow (newBlock, r_{cur})$ 
208:        $curProposal \leftarrow \text{MSG}(\text{PROPOSE}, propData)$ 
209:       broadcast  $curProposal$  once
210:     as a node // Also executed by the leader
211:       start or reset  $T_{blame}(v_{cur})$  to time  $4\Delta$ 
212:       wait for first valid  $newProposal$  from any node
213:       // Vote in the head
214:       update  $B_{lck} \leftarrow newProposal.data[0]$ 
215:       broadcast  $newProposal$  once
216:       set  $T_{commit}(B_{lck}) \leftarrow 4\Delta$ 
217:       call NEXTRound
218:     ▶ View-Change Phase (Safe + Live)
219:     // Handle lack of progress
220:     on time out of  $T_{blame}(v_{cur})$  // Crash
221:        $blameData \leftarrow \perp$ 
222:        $blameMsg \leftarrow \text{MSG}(\text{BLAME}, blameData)$ 
223:       broadcast  $blameMsg$ 
224:     // Handle equivocation
225:     on  $pr_1, pr_2 \wedge pr_1.round = pr_2.round$  // Byz.
226:        $blameData \leftarrow \perp$ 
227:        $blameMsg \leftarrow \text{MSG}(\text{BLAME}, blameData)$ 
228:       broadcast  $(blameMsg, pr_1, pr_2)$  once
229:     on blame  $m$  with  $m.data = (prop_1, prop_2)$  // Byz.
230:       cancel all commit timers  $T_{commit}(\cdot)$ 
231:       broadcast  $m$  once
232:     // Change the view
233:     on  $f + 1$  MATCHINGMSG( $m$ , blame,  $v_{cur}$ ) as  $V$ 
234:       cancel all commit timers  $T_{commit}(\cdot)$ 
235:        $blameQC \leftarrow \text{QC}(V)$ 
236:       broadcast  $blameQC$  once
237:     on valid  $blameQC$ 
238:       broadcast  $blameQC$  once
239:       wait  $\Delta$  // Ensure all correct nodes quit the view
240:       call QUITVIEW
241:   procedure QUITVIEW
242:      $commitData \leftarrow B_{com}$ 
243:      $comReqMsg \leftarrow \text{MSG}(\text{COMMITUPDATE}, commitData)$ 
244:     broadcast  $comReqMsg$ 
245:   wait  $5\Delta$  to obtain  $commitQC$ 
246:   broadcast  $commitQC$  and wait  $\Delta$ 
247:   call NEWVIEW( $commitQC$ )
248:   on COMMITUPDATE msg  $m$  with  $b := m.data$ 
249:     if  $b$  does not conflict with  $B_{lck}$  then
250:       send MSG(CERTIFY,  $b$ ) to the sender of  $m$ 
251:   on  $f + 1$  MATCHINGMSG( $m$ , CERTIFY,  $v_{cur}$ ) as  $V$ 
252:      $commitQC \leftarrow \text{QC}(V)$ 
253:      $commitUpdateMsg \leftarrow \text{MSG}(\text{COMMITUPDATE}, commitQC)$ 
254:   on  $commitQC'$  for  $b'$  from other nodes
255:     if  $b'$  does not conflict with  $B_{lck}$  and extends  $commitQC$  then
256:        $commitQC \leftarrow commitQC'$ 
257:   procedure NEWVIEW( $commitQC$ )
258:      $v_{cur} \leftarrow v_{cur} + 1$ 
259:      $r_{cur} \leftarrow 1$ 
260:      $L \leftarrow \text{LEADER}(v_{cur})$ 
261:     as new leader
262:       wait  $4\Delta$  to hear  $commitQC$  from  $f + 1$  nodes in status
263:        $prop \leftarrow \text{MSG}(\text{NEWVIEWPROPOSAL}, status)$ 
264:       broadcast  $prop$  once // Line 267 calls NEXTRound
265:       wait for  $f + 1$  MATCHINGMSG( $m$ , VOTEMSG,  $v_{cur}$ ) with
266:          $m.data = H(prop)$  as  $V$ 
267:        $qc \leftarrow \text{QC}(V)$ 
268:        $propData \leftarrow qc$ 
269:        $curProposal \leftarrow \text{MSG}(\text{PROPOSE}, propData)$ 
270:       broadcast  $curProposal$  once
271:     as a node
272:       send  $commitQC$  to  $L$ 
273:       set  $T_{blame}(v_{cur})$  to  $8\Delta$ 
274:     on valid NEWVIEWPROPOSAL msg  $\wedge r_{cur} = 1$ 
275:       broadcast  $prop$  once
276:       if  $prop.data$  extends highest block in  $prop.data.status$  then
277:          $data \leftarrow H(prop)$ 
278:          $voteMsg \leftarrow \text{MSG}(\text{VOTEMSG}, data)$ 
279:         broadcast  $voteMsg$ 
280:          $T_{blame}(v_{cur}) \leftarrow 6\Delta$ 
281:         call NEXTRound
282:       on valid PROPOSE msg  $\wedge r_{cur} = 2 \wedge$  valid QC in  $msg.data$ 
283:         broadcast  $msg$ 
284:         call NEXTRound // Go to steady-state
285:   ▶ Commit Rule // Anytime
286:   on time out of  $T_{commit}(block)$ 
287:     update  $B_{com}$  to  $block$ 
288:     commit block and its parents

```

blame messages and also cancel their commit-timers (Line 225) to preserve safety.

Quitting the View. On receiving $f + 1$ blame messages for v_{cur} (Line 227), the nodes build a certificate and broadcast it, signaling other correct nodes to quit the view v_{cur} (Line 231). After sending the blame certificate, correct nodes wait for Δ to ensure that all correct nodes quit the old view v .

After quitting the view v_{cur} , the nodes first obtain certificates for their highest committed block B_{com} . They do this by broadcasting their B_{com} . All correct nodes vote for other's B_{com} if it does not

conflict with the node's local B_{com} and B_{lck} . As we show in the proofs, this will always be true for all correct nodes.

The nodes then broadcast their highest certificate to all the nodes. The other node's update their local highest certificate if the newly received certificates are longer and does not conflict with their local B_{lck} . This ensures that no matter whose certificate the (potentially Byzantine) leader picks along with f other Byzantine server's certificates, it must extend the highest committed block of all the correct nodes. The final highest certificate is input to the next function **NEWVIEW**.

In this step, the nodes wait 5Δ time before starting the next view to ensure that all correct nodes have sufficient time to collect a certificate for their B_{com} . Say, the first correct node quits the view v_{cur} at time t . By time $t + \Delta$, all correct nodes will quit the view v_{cur} . All correct nodes will have broadcast their B_{com} by this time, which will reach all correct nodes by time $t + 2\Delta$, which will have been chain synchronized by time $t + 4\Delta$. By time $t + 5\Delta$, all correct nodes will have collected $f + 1$ votes for their B_{com} .

Starting the New View. In this function, all correct nodes send their certificates for B_{com} to the new leader. The new leader collects certificates from all the nodes and proposes a block containing $f + 1$ such certificates.⁵ The nodes vote for this proposal if it extends the highest block among all the certificates.

A correct leader can always build a valid proposal. Say, the leader is the first node to enter the new view at time t . By time $t + \Delta$, all correct nodes will have entered the new view and sent their certificates. These reach the leader by time $t + 2\Delta$, and accounting for chain synchronization, the leader can construct a proposal by time $t + 4\Delta$. Then it will propose a `NEWVIEWPROPOSAL` message for the round 1 in the new view. It then builds a quorum certificate and proposes the round 2 message, after which it transitions to the steady state.

While the new leader is collecting the highest certificates from all the other correct nodes, we need to ensure that the other correct nodes wait for sufficient time so that they do not blame an correct leader. Otherwise, it can result in liveness problems, as the system changes views all the time. Let the earliest node enter the new view $v + 1$ at time t . The latest correct node (possibly the leader) will enter the view $v + 1$ by time $t + \Delta$ (Lines 233 and 239). The leader will receive the latest locked blocks from all the correct nodes by time $T + 5\Delta$ and it will propose a block at this time which will reach all the correct nodes by time $T + 6\Delta$. Adding 2Δ time for chain synchronization, if the leader is correct, it suffices to wait for 8Δ time in Line 266. By time $T + 9\Delta$, the leader can collect votes from all correct nodes and propose. This proposal will reach all the correct nodes by time $T + 10\Delta$ and time $T + 12\Delta$ after synchronization. However, some correct nodes may have started their timer for the proposal in round 2 of the new view at time $T + 6\Delta$. Therefore, for the second round, a timer of 6Δ will suffice in Line 273 to prevent blaming a correct leader.

If the new leader is Byzantine, as long as it extends one correct node's commit certificate, safety is guaranteed. The intuition is that this certificate is at least the highest committed block from the previous view(s). If the new leader does not make progress in round 1 or fails to produce a valid quorum certificate in round 2, then the timers are triggered, and another view-change protocol begins. This preserves liveness in the protocol.

The drawback of the savings during the steady-state is that during the view-change, all the nodes need to generate certificates for their committed blocks. In this step, the nodes convert their votes in the head to explicit votes.

⁵This step can be optimized to reduce the size, if there are $f + 1$ certificates for the same block.

3.5 Discussion

In this section, we present optimizations, analysis, and general discussions about protocol and model.

Equivocation scenario speedups. We can speed up the view change in the event of an equivocation by skipping the wait to construct a quorum certificate. This follows from the fact that all correct nodes will vote for the equivocation as the signatures must match for all the correct nodes. We can thus save communication in this phase.

Add commands in rounds 1 and 2. For ease of exposition, client commands `CMDs` are not included in rounds 1 and 2 of the view change. This can be added to improve throughput slightly.

Taking advantage of control planes. For instance, CPS nodes may be managed by online trusted control nodes and if such a trusted node is available, we can use it during the view-change to simplify and improve energy-efficiency. Most SMR protocols can drop liveness and only preserve safety, i.e., the nodes will ensure that if two correct nodes output blocks B and B' at height h , then $B = B'$. By forgoing liveness, i.e., if two correct nodes output two blocks B and B' at height h then $B = B'$; this does not state that all correct nodes will output B at height h , when the leader is correct, we can guarantee safety. When the leader equivocates or does not make progress, the nodes can simply stop and wait for the control nodes to change the view. This can improve energy-efficiency.

Security Analysis. We defer the security analysis to the full version [16].

Note on optimizations and security. It is easy to show that these optimizations do not violate the safety and liveness of the system. We can easily extend the proofs provided in [16] to show that the protocols with the optimizations are still secure.

4 ENERGY ANALYSIS

In this section, we formalize the energy-measurement model, argue for the need to optimize the common-case and push complexities to the worst-case. We also develop a framework to compare protocols that allows us to make better decisions.

The need for best-case optimality. Let \vec{X} define a vector, which consists of system parameters such as $\vec{X} = (n, f, m, S, R, \sigma_s, \sigma_v)^T$, where n is the total number of nodes tolerating f faults; m is the maximum supported payload size, e.g., the size of `CMDs`; S and R are the costs to send and receive per byte; and σ_s and σ_v are the costs to sign and verify digital signatures.

Let ψ denote the energy cost function of a leader based protocol per unit of consensus (typically a block of client requests). The function takes a column vector \vec{X} of energy-costs of primitives and computes the energy cost of the protocol. An example is: $\psi(\vec{X}) = c_1 m + c_2 n + c_3 m n + c_4 m n S + c_5 m n R + c_6 \sigma_s + c_7 n \sigma_v$ for some constants $c_1, c_2, c_3, c_4, c_5, c_6$ and c_7 . This function ignores the common state machine replication costs such as verifying the semantic validity of client requests, checking if the client has correctly generated the requests, execution of a request, etc. When comparing two different protocols, we drop \vec{X} , and just use ψ . In later section, we use such functions to model protocols and perform energy analysis.

We denote a particular protocol by placing its name in the superscript (e.g., $\psi^{Baseline}$), the best-case energy-cost (without faults) of the protocol using subscript B such as ψ_B , the worst-case (with

faulty leader and faults)⁶ energy cost and the view change energy costs (costs to change the leader) using the subscripts W and V respectively such as ψ_W and ψ_V . We assume $\psi_V = \psi_W - \psi_B$ ⁷⁸.

We wish to design a protocol ψ that is more energy efficient than another protocol ψ^* . Trivially, if $\psi_B \leq \psi_B^*$ and $\psi_W \leq \psi_W^*$, then we have already achieved our goal. Therefore, let us consider the case where $\psi \leq \psi^*$ is not true always. Let N be the total number of blocks agreed upon in which W worst-case events occur which result in $V = W$ view changes. Then using $(N - V) \psi_B + V (\psi_W) \leq (N - V) \psi_B^* + V (\psi_W^*)$, we get $v_f = \frac{V}{N} \leq \frac{\psi_B^* - \psi_B}{\psi_W - \psi_V}$. Let $v_f = \frac{V}{N}$ be the ratio of view changes V to total number of blocks N . The fraction of *good* (best-case) runs are $1 - v_f$ with $0 \leq v_f \leq 1$.

(Un)Favorable conditions. There are two regions of solutions for the above inequality when $v_f \geq 0$: (i) $\psi_B > \psi_B^*$ (worse best-case) with $\psi_V < \psi_V^*$ (better worst-case), which we term as the *worst-case optimal* solution set, and (ii) $\psi_B < \psi_B^*$ (better best-case) with $\psi_V > \psi_V^*$ (worse worst-case), which we term as the *best-case optimal* solution set.

This means that we want to either make the protocol worst-case efficient (region i) or best-case efficient (region ii) for ψ to be more energy efficient than ψ^* . In bounded-synchrony, since the number of faults are bounded by f , V is also bounded by f . Therefore, for our setting, asymptotically we need a best-case optimal ψ , i.e., the *best-case optimal* solution set. We also want $\psi_V - \psi_V^*$ to be as close as possible. In other words, if ψ_V is much better than ψ_V^* , then the number of best-case events $N - V$ we need to run the protocol will be large, in which case a worst-case optimal solution set may be better.

$$N \geq V \left(\frac{\psi_V - \psi_V^*}{\psi_B^* - \psi_B} \right)$$

For non-synchronous systems, observe that the adversary can choose which phase the protocol spends the most time on. As a result, an adversary with the goal of thwarting energy-efficiency can always do so by ensuring that the protocol is spending most of its time in the non-optimal phases. Hence, we restrict ourselves to bounded-synchronous systems in this work and leave the optimization problem for non-synchronous systems an interesting open problem for the future.

For any network setting, Dolev and Strong [32] showed in Theorem 4.1 that in the worst-case it is impossible to avoid the $f + 1$ round lower bound. For our synchronous setting, the network is always reliable within the Δ parameter. In a leader based protocol we will have at most f Byzantine leaders. Thus, intuitively, protocols must be designed to be best-case optimal.

⁶By worst case, we ignore a denial-of-service adversary that spams invalid messages. Such adversaries are out-of-scope of this work and addressing them is an orthogonal problem [4, 6, 8, 22]. For example Elmany et al. [34] classify DoS and jamming as a network-layer concern and not application layer (where we design protocols).

⁷Note that we do not assume that $\psi_V > 0$. There exist protocols where the protocol aborts early in case of Byzantine failures such as Tendermint [20], which results in $\psi_V < 0$.

⁸What about cases that are neither best or worst-case? Examples of such cases are nodes that send incorrect messages, spam messages, etc. We ignore such cases as they are an orthogonal problem common to all SMR protocols and can be resolved by heuristics such as reputation or banning mechanisms for such peers.

THEOREM 4.1 (LOWER BOUND ON BA [32]). *Byzantine agreement using authentication can be achieved for n processors with at most f faults within $f + 1$ phases, assuming $n > f + 1$.*

This bound reiterates that we need to make protocols energy efficient in the best-case scenario with lower view-change costs to maximize f_e , compared to the baseline justifying the fit of EESMR when compared to other protocols.

In practice, protocols may exist where a view-change may exceed the energy available (say the energy available in a fully-charged battery). In such conditions, the worst-case optimal solution might be more practical, or energy-efficiency may not be possible for the given f faults and the given energy-budget.

Adversary for Energy-Efficiency. Let f_s be Byzantine nodes that attack safety. These may send incorrect messages, but send messages when they must, i.e., they do not attack liveness. Let f_c be Byzantine nodes that attack liveness (they do not send messages, but if they do, they will send the correct messages). For safety, we have $n > f_s$, and for liveness we have a necessary condition $n > 2f_c$.

We explore a similar relationship for energy faults f_e , i.e., faults that wish to drain the energy of correct nodes. If we want a protocol to be energy fault-tolerant, then the number of faults $f_e \leftarrow f$ should be such that f_e of the worst-cases for ψ should still be better than f_e of the worst-cases for ψ^* . Solutions for f that satisfy $\psi(f, \cdot) \leq \psi^*(f, \cdot)$ dictate the energy bound. This is true for example, when ψ and ψ^* are monotonic in f . For our setting, let f_e Byzantine nodes cause the worst case scenario to occur f_e times after which finally a best-case scenario occurs. Then using $f_e \cdot \psi_W^{EESMR} + \psi_B^{EESMR} \leq \psi^{Baseline}$, we obtain

$$f_e \leq \frac{\psi^{Baseline} - \psi_B^{EESMR}}{(\psi_B^{EESMR} + \psi_V^{EESMR})} \quad (EB)$$

In many applications, despite the main protocol tolerating up to 50% faults, the actual protocol might tolerate $f < f_{max}$ due to connectivity bottlenecks or other reasons. Obtaining f_e in relation to another protocol helps to choose the more energy-efficient protocol.

Our analysis allows administrators and protocol stake-holders and deployers to model protocols and use the application details such as expected number of faults, communication and computation primitives used to make energy-aware protocol choices.

5 CPS AND PERFORMANCE ANALYSIS

In this section, we consider the Cyber-Physical System (CPS) context to evaluate the energy efficiency of our protocol. Our choice of CPS is motivated by two different observations. a) The energy-starved nature of conventional CPS nodes makes it essential to ensure that the energy overhead of a BFT protocol is minimal on the system. b) Conventional machines run a heavy operating system with millions of CPU instructions. Calculating the energy consumed by the protocol from the machine's overall energy consumption requires us to separate the operating system's inherent energy consumption, which is prone to a lot of noise. This noise masks the energy consumed by the BFT protocol. Hence, we consider low-powered CPS nodes, which run a lightweight OS with much lower energy consumption. Consequently, the noise contributed by the




Figure 1: An example feasible region plot for different message sizes (m) and number of nodes (n). This plot is generated for RSA-1024, and the nodes communicate with each other using Wi-Fi and the external trusted node uses 4G. BL is Baseline. The z-axis is the difference between EESMR energy cost and the baseline. When negative, EESMR is more energy-efficient.

OS to the energy readings of the device is also lower, allowing us to measure the energy of the protocol accurately.

5.1 Protocol Energy Comparison

In the previous sections, we designed our EESMR to be more energy-efficient than the state-of-the-art protocol Sync HotStuff. However, every application is unique and may contain several communication media along with trusted control nodes. EESMR may not be the most optimal solution in all of these scenarios. In this part, we present analysis techniques to compare energy-efficiency of protocols.

Comparison with trusted-baseline. We compare our work with a trusted-baseline protocol using theoretical tools. In this baseline protocol, we assume the existence of a trusted node. Trivially, if the trusted node uses the same communication network as the other nodes, then the trusted-baseline is always energy-efficient. We consider the case where the trusted node is on a different (more expensive) communication medium than the communication medium among nodes (e.g., satellites to communicate with the trusted node and ground-stations for local communication among other nodes). An example of such scenario is communication between CPS/IoT devices and the control servers.

The baseline protocol assumes that all the CPS nodes are directly connected to the trusted node using the expensive medium and not use the links between the CPS nodes. The baseline still assumes that f out of n CPS nodes are Byzantine.

We analytically count operations and build equations for $\psi_{B}^{Baseline}$, ψ_{B}^{EESMR} and ψ_{V}^{EESMR} in MATLAB software. We use the energy measurement values observed by related works and from our experiments in Section 5.4 in our analysis. We present an example

feasible region in Fig. 1 which shows when EESMR is favorable over a trusted-baseline protocol (BL) assuming an external trusted party over a slightly more expensive medium 4G, while the CPS nodes use Wi-Fi to run EESMR. Such modeling is useful for example, to choose the more efficient protocol for a given set of parameters. This analysis can be extended to consider our crash-fault tolerant version, or our steady-state can be combined with the trusted-baseline only for the view-change to improve the energy-efficiency.

5.2 CPS System Model

Multicasts. We take advantage of the multicasts available in the CPS setting by modeling the network as hypergraphs ([16]). The wireless multicast for instance guarantees that all correct nodes hear a consistent value in a time-bounded fashion (see note on jamming). We assume that all nodes have access to a multicast channel with at least k nodes such that the system remains f -connected, i.e., every node can connect with $f + 1$ nodes.

Notations. We refer to these multicast channels as k -casts as we can multicast once to reach k nodes. These are realized using Ethernet multicasts or wireless multicasts in the CPS system. Our protocols are updated accordingly, to use these k -casts instead of broadcast, and an appropriate Δ parameter is used to ensure that all correct nodes receive a message sent by any other correct node. We denote by D_{in} and D_{out} the number of incoming and outgoing k -cast links. *Note on Bounded-Synchrony for CPS.* The employed equivocation detection is only applicable to bounded-synchronous (CPS) networks, which is viable in closed controlled environments [25, 35, 42, 59, 73, 74, 80] such as remote, industry and military applications. Notice that for asynchronous and partially synchronous network models, due to the unknown delays, an adversary can always deplete any finite energy resources of the CPS devices, making any meaningful energy analysis challenging. The energy-efficiency analysis of such protocols during periods of synchrony is of independent interest and presents an interesting future work.

Note on Jamming. We assume that the hyper-edges are synchronous [25, 31, 35, 42, 59, 71–74, 80], and similar to the related works [4, 6, 22, 56–58, 60, 63] consider network jamming to be an orthogonal problem. Nevertheless, we note that the location of deployment such as data-centers, military networks using satellites, or remote large areas will dictate whether jamming is a part of the threat model. If large-scale jamming is a part of the threat model for a setting, unicast links⁹ are reliable. The use of unicasts do not affect our claims on energy-efficiency over prior SMR protocols.

We measure and compare the costs of multicasts and unicasts and measure their reliability. We implement our proposed protocol EESMR and the state-of-the-art protocol Sync HotStuff [6] on CPS devices and compare their energy costs.

5.3 System setup

We implemented our protocol and the baseline protocol Sync HotStuff in C++ and tested them on NUCLEO-F401RE nodes with ARM Cortex M4 84 MHz processors. The main board has 512 kB of flash memory and 96 kB of SRAM. The nodes communicate using Bluetooth Low Energy (BLE) modules. We used Saleae Logic-Pro 8

⁹Jamming in unicast links in wireless networks can be mitigated by Spread-Spectrum Frequency Hopping [69].

and INA169 sensors to measure the energy consumed by the nodes running the protocols.

5.4 Evaluating primitives

We measure the energy costs of the communication and cryptographic primitives described in Section 2.

Communication primitives. We measured the cost of sending a message through different communication media in Table 1. The energy reading for BLE is the energy needed to transmit one packet of data. However, this BLE transmission does not verify if the other party received it reliably. We notice that BLE’s energy requirement is two orders of magnitude lower than WiFi and three orders of magnitude lower than 4G, which is the primary reason of us using BLE in our implementation. We evaluated this mode of communication further by measuring the energy costs of sending and receiving messages of various sizes using unicasts and multicasts in BLE. We use BLE advertisement packets as multicasts and k -casts. We observe that these packets have payloads limited to 25 bytes per packet by the BLE GAP specification. For large messages, we must therefore fragment them into multiple advertisements. BLE advertisements are packets in the link layer, with no inherent capability to handle packet losses. We tackled this inherent unreliability using redundant transmissions. We conducted experiments to chart the reliability of k -casts. The setup involved a set of BLE-enabled embedded devices placed within a 10 meter distance of each other, which is within the range of BLE. We transmitted batches of 10,000 packets with varied levels of redundancy and measured the number of successful packets received by the receivers. The failure rate was measured against the redundancy factor of transmission and the energy values were measured accordingly. A k -cast is considered to be a success only if all the k receivers successfully receive the data. We repeated this experiment for different values of k and recorded the experimental results in Fig. 2a. We observe that the failure rates exponentially decrease with the increase in the redundancy factor and energy spent per transmission. The probability of a transmission failure increases with the value of k and the energy needed to guarantee a 99.99% reliability also increases accordingly. However, for specific applications, a threshold of 99.99% might not be sufficient. For such applications, the reliability can be customly set using concepts like erasure and error-correcting codes, where the amount of redundancy can be set according to the desired reliability. Based on the noise in the environment, we can measure a value Δ such that the desired reliability is achieved.¹⁰

With the added redundancy, we experimentally observe that it takes bounded 200 ms to transmit a 25 byte message with 99.99% reliability over a multicast link in BLE, with $k = 7$. BLE has a low bandwidth resulting in a high transmission time, but it is energy efficient when compared to other alternatives. The experiments were conducted in a laboratory with many rogue advertisement packets from other devices which increased the energy consumption of the receiver nodes. But this also mimics a practical setting due to the ubiquity of BLE. For all further experiments, we use this setting of k -casts that guarantees 99.99% reliability for each k -cast link.

¹⁰It can be an interesting optimization to improve the reliability without redundant transmissions using erasure coding. However, we note that erasure coding solutions typically employ large pre-computed tables, which will not fit on the employed CPS boards without non-trivial implementation efforts.

Multicasts vs. unicasts on BLE. We explore the unicast alternative to multicasts. Unicasts offers better reliability at the cost of more point-to-point transmissions. The unicast links use the GATT (Generic Attribute Profile) protocol, which is a connection-based protocol in the BLE realm. GATT offers reliable data transmission by inherently handling packet drops and re-transmitting data when necessary. Fig. 2b compares the energy consumption of 99.99% k -cast on BLE with the equivalent $D_{out} = k$ unicast links for BLE to transmit messages with different payloads. We observe that expectedly, the energy required to transmit a k -cast using equivalent unicasts increases linearly with the k .

Fig. 2b compares the energy costs to transmit the same payloads using the 99.99% reliable multicast and equivalent unicasts. We observe that a unicast link is more effective compared to a multicast for bigger payloads, but this advantage is quickly negated as the value of k increases. The embedded devices used in our experiments cannot handle concurrent unicast connections with its neighbors, due to which data transfers using unicasts add extra time overheads.

5.5 Cryptographic Primitives

Hash and HMAC costs. We instantiate the MAC (message authentication code) algorithm using SHA-256 as the underlying hash function, and measure the hashing cost for different message sizes. We use short keys of recommended size 64 bytes. The cost of signing (μ_S) and verifying (μ_V) is the same as for the HMAC scheme. The major cost in the HMAC scheme was mostly due to the underlying SHA-256 algorithm. We found that the cost of hashing increased linearly with message size.

Public key primitives. We measured the energy costs for Elliptic Curve Digital Signature Algorithm (ECDSA) [43] and RSA [28] for various security parameters, and recorded their energy readings in Table 2. For ECDSA [43], we found that brainpool curves (BP-XXX) [51] were generally expensive to sign (5J), and verify (11J) for 160 bit curves. NIST optimized curves (SECP-XXX) [36] offer better performance in comparison (1J and 2J respectively) for 160-bit curves. We use the implementations from MbedTLS [3] (commit: b6229e304) for our measurements. We implemented BP160R1 using MbedTLS and RFC 5639 [51]. Finally, RSA using 1024-bit modulus costs 400mJ and 20mJ to sign and verify, being an ideal candidate for CPS.

RSA using 1024-bit modulus provides 80-bits of security and should be practical for most CPS settings [2, 46]. RSA provides two benefits over ECDSA: (i) reduced energy costs, and (ii) the benefit of asymmetry in verification and signing costs. The latter implies higher energy costs for Byzantine nodes (to equivocate) while the impact is significantly less on the correct nodes that verify.







5.6 Evaluating EESMR

We implemented EESMR on embedded devices using BLE advertisements as k -casts with sufficient redundancy to achieve a reliability of 99.99% for message transmission. Considering the memory limitations of the devices we used, we performed experiments on the blocking variant of EESMR.¹¹ We use $D_{in} = k$, and $D_{out} = 1$ in our

¹¹The non-blocking variant of our protocol requires more memory (theoretically unbounded) to process all the proposed blocks. The only changes required in the protocol are to let the timers run concurrently, and multiple proposals can be forwarded by a node. The energy analysis still hold with respect to every block. A common issue

Table 1: Sample energy consumption data for different media. All the measurements are in milliJoule (mJ).

Message Size	BLE			4G LTE [41]		WiFi	
	Send	Recv	Multicast	Send	Recv	Send	Recv
256 B	0.73	0.55	0.58	494.84	69.54	81.2	66.66
512 B	1.31	1.11	1.17	989.68	139.08	153.98	123.23
1 kB	2.93	2.64	2.35	1979.36	278.17	310.54	231.52
2 kB	5.91	5.23	4.70	3958.72	556.35	610.55	423.58

**(a)** The failure rate of k -casts as a % mapped against the energy consumed by the sender and receiver. We achieve a reliability rate of 99.99% by spending 5.3 mJ per message by the sender and 9.98 mJ per message for the receiver.**(b)** Comparing costs of unicast vs. multicasts. We use a reliability of 99.99% for k -casts. UC stands for Unicast, S stands for Sender and R means receiver.**(c)** Average energy per state machine replication consumed by a correct EESMR leader and other nodes and its variation with k . We use $|b_i| = 16$ bytes.**(d)** Energy consumed by a correct EESMR leader per SMR for variable block sizes with respect to the value k .**(e)** Energy consumed by EESMR leader per view change operation for different f , evaluated with $n = 15$ nodes. We use $|b_i| = 16$ bytes.**(f)** Total energy consumed by the correct nodes in EESMR and Sync HotStuff with respect to the total number of nodes n .**Figure 2: Energy characterization of various scenarios realized on embedded devices****Table 2: Comparing energy costs (in J) for signature generation (σ_s) and verification (σ_v) for ECDSA curves and RSA schemes.**

Algorithm	Parameters	Sign (in J)	Verify (in J)
ECDSA	BP160R1	5.80	11.03
	BP256R1	13.88	27.34
	SECP192R1	0.84	1.50
	SECP192K1	1.16	2.24
	SECP224R1	1.10	2.14
	SECP256R1	1.60	3.04
	SECP256K1	1.72	3.35
RSA	1024-bit modulus	0.40	0.02
	1260-bit modulus	0.79	0.03
	2048-bit modulus	2.41	0.06
HMAC	–	0.19	0.19

experiments, where each hyper-edge has degree k . The network topology is defined as follows: In a system of $N = \{p_1, \dots, p_n\}$ nodes using k -casts, every node p_i transmits messages to nodes $p_{i+1 \bmod n}, \dots, p_{i+k \bmod n}$. Every node p_i receives messages from $p_{i-1 \bmod n}, \dots, p_{i-k \bmod n}$, because of which only these nodes need to be in communication range. The physical placement of the nodes is such that each node is within communication range of all nodes with which it has an edge. We measured the energy consumption of EESMR by varying n , k and the block size. We take the Δ value to be $10n$ seconds. We took this metric because of the inability of our devices to scan and transmit data simultaneously. We designed a time-sequenced schedule and allocated an interval of 10 seconds

with non-blocking protocols (all works so far) is that an adversary can propose a large number of non-blocking proposals but revert them during a view-change leading to wasted energy. We leave this as an orthogonal problem of independent research, as it is generally applicable to all non-blocking protocols. This can be mitigated by bounding the number of allowed non-blocking proposals.

for every node to transmit data reliably. In this interval, all other nodes will only listen to packets from the designated node scheduled for transmission. Even if Byzantine nodes try and transmit their messages in undesigned intervals, the honest nodes will ignore their packets and only listen to the designated node. We show that the network is synchronous with this value of Δ . However, in rare situations where the network's Δ is higher, honest nodes can trigger the view change protocol to replace the leader, even when the leader was honest. We instantiate the abstract block structure, concretely as $B = \langle m, H(b_m), H(h_{m-1}), \langle i, H(b_i) \rangle_L \rangle$, where H is a cryptographic hash function, m is the block height, and L is the leader. In CPS, these blocks contain data b_m to maintain state.

We also implemented optimizations for the equivocation scenario, by avoiding construction of quorum certificates and using the equivocation and proof to quit the view. We also optimize the no-progress scenario for the blocking version of our protocol, by avoiding the highest certificate construction, and ensuring that the status message only contains node's locked block. This is secure because the highest committed block is either the locked block or its parent, and thus it is safe to extend any one of the locked blocks among $f + 1$ locked blocks.

We recorded the energy consumed by the protocol by subtracting the energy consumed by the device in sleep state. While doing so, we noticed that each node consumes 0.3 mW of power while in the sleep state and consumes approximately 1 mW of power while conducting SMR. Even at a modest frequency of one SMR per hour, the energy consumed by the SMR causes a substantial load on the node's limited energy resources. Hence, energy-optimal SMR protocols are necessary to reduce the overhead of SMR on lightweight nodes.

Our first observation is that the energy cost of EESMR is independent of n in the best case as we do not use certificates (which are $f + 1$ signatures). The energy cost only depends on k . We show this in Fig. 2f, where the total energy consumption of all the correct nodes increases linearly with respect to the number of correct nodes in the network.

Next, we measured the difference in energy costs between the leader and the other correct nodes for different values of k . In Fig. 2c, we observed this difference. The linearity of energy consumption with k is due to k incoming edges. The two energy costs are close, with the leader having a higher cost. We also evaluated the variation of energy consumption with block size and show that EESMR scales well with increasing message payloads (Fig. 2d).

We measured the energy consumed by EESMR to perform a view change operation, in case of a Byzantine leader. The energy consumed per view change is plotted (Fig. 2e) with respect to the value of f , in case of an equivocating leader and a stalling leader. The value of k has been taken to be $f + 1$ to measure the minimum energy that a leader spends in the view change. Comparison with correct case SMR has also been performed with $k = f + 1$. The stalling case is much more expensive in terms of energy because of the vote casting and the $f + 1$ blame messages that need to be verified. It should be noted that the energy measurement for the view change is only for the blocking version of EESMR.




Figure 3: The energy consumed by a leader in EESMR vs. Sync HotStuff to tolerate f Byzantine faults in a system of $n = 13$ nodes. We show both the faulty and non-faulty leader cases.

5.7 Comparison with Sync HotStuff

We compare in Fig. 3 EESMR's energy readings with our implementation of Sync HotStuff [4]. We made simplifying assumptions in favor of Sync HotStuff, by partially implementing vote forwarding. The minimum energy to tolerate f faults has been measured for EESMR and Sync HotStuff. We designed our $n = 13$ graph to ensure that it is f fault-tolerant. We measure the energy consumed by a correct leader in a network consisting of $\lfloor \frac{n}{2} \rfloor$ faulty nodes. The amount of energy consumed by a node linearly increases with the value of k . As k increases, every node receives messages from an increasing number of nodes, which increases the energy spent in receiving them. Additionally, Sync HotStuff uses a certificate of size $f + 1$ increasing the energy costs. When the leader is correct, Sync HotStuff is 2.85× more energy hungry than EESMR. In a view change, the ratio of the energy consumed by EESMR Energy with respect to Sync HotStuff is 2.05. To repeat our earlier argument, in the steady state, the number of SMR rounds where there will be a view change will become far smaller than where the leader will be benign. We also compare the total energy consumed by correct nodes per SMR for both protocols in Fig. 2f. This demonstrates the scalability of EESMR in the best-case.

6 RELATED WORK

To the best of our knowledge, there is no prior work that formally analyzes SMR protocols for energy-efficiency.

SMR. Prior works [6, 8, 15, 20, 21, 39, 48, 76] focused on improving the message complexity or the commit latencies of SMR. But nodes in CPS tend to be constrained in their energy and computational resources, and decreasing communication complexity and increasing the computation does not help the CPS nodes. Protocol designs as of today use votes and cryptographic certificates while completely ignoring their energy costs on CPS devices. In the absence of signature aggregation techniques, these protocols [6, 39, 76] increase the energy requirements for SMR. We present a comparison of the related works in Table 3.

Let δ be the actual network speed between any two nodes in a fully connected system. OptSync [63] commits and produces a block every 2δ . To achieve this all nodes verify $3n/4 + 1$ signatures. Sync HotStuff [6] on the other hand produces a block every 2δ , has a commit latency of 2Δ , and verifies $n/2 + 1$ signatures. Clearly, Sync

Table 3: Comparison of best-case and worst-case scenarios in related SMR works. Here, the number of nodes is n ; We consider a partially connected network with all nodes connected to d neighbors such that the graph remains connected even if f nodes are Byzantine. Some of the protocols are for fully connected graphs but we compare them in our partially connected setting. Block period refers to the time difference between two successive proposed blocks. $\delta \leq \Delta$ represents the actual message delivery time.

Protocol	Correct Leader (best-case)				Faulty Leader (worst-case)			
	Communication Complexity	Public Key Ops.		Block Period	Communication Complexity	Public Key Ops.		Block Period
		Sign	Verify			Sign	Verify	
Abraham <i>et al.</i> [4]	$O(n^2d)$	$O(n)$	$O(n^2)$	—	$O(n^3d)^*$	$O(n)$	$O(n^2)$	—
Sync HotStuff [6]	$O(n^2d)$	$O(n)$	$O(n^2)$	2δ	$O(n^3d)^*$	$O(n)$	$O(n^2)$	14Δ
OptSync [63]	$O(n^2d)$	$O(n)$	$O(n^2)$	2δ	$O(n^3d)^*$	$O(n)$	$O(n^2)$	14Δ
Rotating BFT SMR [8]	$O(n^2d)$	$O(n)$	$O(n^2)$	2δ	$O(n^2d)^\dagger$	$O(n)$	$O(n^2)$	14Δ
EESMR	$O(nd)$	$O(1)$	$O(n)$	0	$O(n^3d)^*$	$O(n)$	$O(n^2)$	21Δ

*A Byzantine leader can make the communication complexity in a view change to be $O(n^3d)$, but this results in all correct nodes committing $O(n)$ blocks after the view change. † For the rotating leader protocol, we report the worst-case for one iteration with a bad leader. In theory, we could have $o(n)$ bad leaders consecutively, leading to an increase in all the numbers by a factor of n .

HotStuff is more energy efficient than OptSync with an increase in latency from 2δ to 2Δ . EESMR can create blocks as soon as it has enough messages without any other delays giving it a block period of 0. It has a commit latency of 4Δ (accounting for chain-synchronization), and uses only 1 signature. Thus, EESMR is more efficient than both Sync HotStuff and OptSync to commit blocks.

Efficient SMR. MinBFT [70] and CheapBFT [44] are partially synchronous SMR protocols that use trusted counters and subsystems to improve the amount of communication of PBFT [21]. They use trusted components to assign sequence numbers to requests, thereby restricting the equivocating capabilities of Byzantine adversaries. They also use only $f + 1$ active replicas which replicate the state thereby saving system resources. Apart from the stronger assumptions, they still use certificates per block and 3 rounds of all-to-all diffusion, leading to the similar energy inefficiencies.

Other metrics of efficiency such as bit-complexity [12, 29, 33, 55] or latency [5, 7, 8] have been considered in the literature. However, they use impractical techniques such as expander graphs or computationally expensive primitives such as threshold cryptography or online error correction.

There is an orthogonal set of works that improve the efficiency of protocols by using trusted components [30, 75]. Indeed, as evidenced in Fig. 1, trusted components can significantly lower energy costs. But existing works assume all-to-all communication which makes the protocol flood $O(n)$ messages in every round. However, it remains an interesting open problem to combine our techniques with trusted components for improved energy-efficiency.

Using multicasts. Multicasts occur frequently due to the spatial arrangement of CPS nodes and the omni-directional nature of wireless communication media. It may not cost more energy to use multicast links over unicast links. Several prior works [27, 37, 40, 45, 47, 53] leverage multicasts to address BA. A common drawback of these works is that they are not tailored to leverage practical multicasts, e.g., considering the drop in reliability with increasing degree.

Koo *et al.* [47] provide a bound on the number of Byzantine nodes that can be tolerated in the neighborhood of a correct sender while using multicasts to achieve BA. Montemanni *et al.* [53] and Guo *et al.* [40] propose models in determining the optimal transmission

power of a sender. However, these works [40, 47, 53] cannot be directly used for CPS since the transmission power of the antennae is typically fixed.

Fitzi *et al.* [37] assume the existence of reliable multicasts between every three nodes. They show that the bound on the number of faulty nodes that can be supported in such a fully connected synchronous network can be improved from $n/3$ to $n/2$ to solve BA. Considine *et al.* [27] show that a resilience of $\frac{n}{f} > (k+1)/(k-1)$ can be achieved in a system with n nodes and k -casts, with f of them being faulty. Khan *et al.* [45] consider multicasts as undirected edges in a graph and provide necessary and sufficient conditions for BA in their model. The main drawback in these works is that there is an assumption that every subset of k nodes have a reliable k -cast present. In this work, we make weaker multicast assumptions of links existing only between neighboring nodes.

CONCLUSION

We present EESMR, an SMR protocol with improved communication in partially connected networks, signature generation, and verification costs. We present analysis techniques to determine choice of protocols and argue why optimizing the best-case is important. We provide a general hypergraph model that can take advantage of multicasts in wireless CPS, when available. Finally, we empirically show a 33 – 64% reduction in energy costs in the steady-state as compared to the state-of-the-art solution Sync HotStuff.

7 ACKNOWLEDGMENTS

We would like to thank the various reviewers and our shepherd Jérémie Decouchant for the helpful feedback and suggestions to improve this draft significantly. We would like to thank Samarjit Chakraborty, Kartik Nayak and Nibesh Shrestha for helpful feedback about our draft. This work was supported in part by NIFA award number 2021-67021-34252, the National Science Foundation (NSF) under grant CNS1846316 and National Science Foundation Cyber Physical Systems (CPS), the United States Department of Agriculture, and the Army Research Lab Contract number W911NF-2020-221.

REFERENCES

- [1] 2020. A Byzantine Failure in the Real World. <http://blog.cloudflare.com/a-byzantine-failure-in-the-real-world/>
- [2] 2021. IBM Documentation. <https://prod.ibmdocs-production-dal-6099123ce774e592a519d7c33db8265e-0000.us-south.containers.appdomain.cloud/docs/en/zos/2.3.0?topic=certificates-size-considerations-public-private-keys>
- [3] 2022. README for Mbed TLS. Mbed TLS. <https://github.com/Mbed-TLS/mbedtls>
- [4] Ittai Abraham, Srinivas Devadas, Danny Dolev, Kartik Nayak, and Ling Ren. 2019. Synchronous Byzantine Agreement with Expected $O(1)$ Rounds, Expected $\$O(N^2)\$$ Communication, and Optimal Resilience. In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Goldberg Ian and Tyler and Moore (Eds.). Vol. 11598 LNCS. Springer International Publishing, Cham, 320–334. https://doi.org/10.1007/978-3-030-32101-7_20
- [5] Ittai Abraham and Danny Dolev. 2015. Byzantine Agreement with Optimal Early Stopping, Optimal Resilience and Polynomial Complexity. arXiv:arXiv:1504.02547 <http://arxiv.org/abs/1504.02547> Comment: full version of STOC 2015 abstract.
- [6] Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Maofan Yin. 2020. Sync HotStuff: Simple and Practical Synchronous State Machine Replication. In *2020 IEEE Symposium on Security and Privacy (SP)*, Vol. 2020-May. IEEE, Oakland, 106–118. <https://doi.org/10.1109/SP40000.2020.00044>
- [7] Ittai Abraham, Kartik Nayak, Ling Ren, and Zhuolun Xiang. 2021. Good-Case Latency of Byzantine Broadcast. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. ACM, New York, NY, USA, 331–341. <https://doi.org/10.1145/3465084.3467899> arXiv:2102.07240
- [8] Ittai Abraham, Kartik Nayak, and Nibesh Shrestha. 2022. Optimal Good-Case Latency for Rotating Leader Synchronous BFT. In *25th International Conference on Principles of Distributed Systems (OPODIS 2021) (Leibniz International Proceedings in Informatics (LIPIcs)*, Vol. 217), Quentin Bramas, Vincent Gramoli, and Alessia Milani (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 27:1–27:19. <https://doi.org/10.4230/LIPIcs.OPODIS.2021.27>
- [9] Ittai Abraham, Benny Pinkas, and Avishay Yanai. 2020. Blinder: MPC Based Scalable and Robust Anonymous Committed Broadcast. *Cryptology ePrint Archive* (2020). <https://eprint.iacr.org/2020/248>
- [10] Ittai Abraham and Gilad Stern. 2020. Information Theoretic HotStuff. *arXiv preprint arXiv:2009.12828* (2020). arXiv:2009.12828
- [11] Amitanand S. Aiyer, Lorenzo Alvisi, Rida A. Bazzi, and Allen Clement. 2008. Matrix Signatures: From MACs to Digital Signatures in Distributed Systems. In *Distributed Computing (Lecture Notes in Computer Science)*, Gadi Taubenfeld (Ed.). Springer, Berlin, Heidelberg, 16–31. https://doi.org/10.1007/978-3-540-87779-0_2
- [12] Nicolas Alhaddad, Sisi Duan, Mayank Varia, and Haibin Zhang. 2021. Succinct Erasure Coding Proof Systems. *Cryptology ePrint Archive* (2021). <https://eprint.iacr.org/2021/1500>
- [13] Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, and Julian Loss. 2022. Network-Agnostic Security Comes for Free in DKG and MPC. <https://eprint.iacr.org/2022/1369>
- [14] Zuzana Beerliová-Trubíniová, Martin Hirt, and Micha Riser. 2007. Efficient Byzantine Agreement with Faulty Minority. In *Advances in Cryptology – ASIACRYPT 2007 (Lecture Notes in Computer Science)*, Kaoru Kurosawa (Ed.). Springer, Berlin, Heidelberg, 393–409. https://doi.org/10.1007/978-3-540-76900-2_24
- [15] Alysson Bessani, João Sousa, and Eduardo E.P. Alchieri. 2014. State Machine Replication for the Masses with BFT-SMART. In *Proceedings of the International Conference on Dependable Systems and Networks*. IEEE, 355–362. <https://doi.org/10.1109/DSN.2014.43>
- [16] Adithya Bhat, Akhil Bandarupalli, Manish Nagaraj, Saurabh Bagchi, Aniket Kate, and Micheal K. Reiter. 2023. EESMR: Energy Efficient BFT-SMR for the masses. arXiv:2304.04998 [cs.CR]
- [17] Umesh Bodkhe, Dhruv Mehta, Sudeep Tanwar, Pronaya Bhattacharya, Pradeep Kumar Singh, and Wei-Chiang Hong. 2020. A Survey on Decentralized Consensus Mechanisms for Cyber Physical Systems. *IEEE Access* 8 (2020), 54371–54401. <https://doi.org/10.1109/ACCESS.2020.2981415>
- [18] Dan Boneh and Xavier Boyen. 2008. Short Signatures Without Random Oracles and the SDH Assumption in Bilinear Groups. *Journal of Cryptology* 21, 2 (April 2008), 149–177. <https://doi.org/10.1007/s00145-007-9005-7>
- [19] Gabriel Bracha and Sam Toueg. 1985. Asynchronous Consensus and Broadcast Protocols. *J. ACM* 32, 4 (Oct. 1985), 824–840. <https://doi.org/10.1145/4221.214134>
- [20] Ethan Buchman, Jae Kwon, and Zarko Milosevic. 2019. The Latest Gossip on BFT Consensus. (Nov. 2019), 1–14. arXiv:1807.04938 <http://arxiv.org/abs/1807.04938>
- [21] Miguel Castro and Barbara Liskov. 2002. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Transactions on Computer Systems* 20, 4 (Nov. 2002), 398–461. <https://doi.org/10.1145/571637.571640>
- [22] Hua Chai and Wenbing Zhao. 2014. Byzantine Fault Tolerant Event Stream Processing for Autonomic Computing. In *2014 IEEE 12th International Conference on Dependable, Autonomic and Secure Computing*. 109–114. <https://doi.org/10.1109/DASC.2014.28>
- [23] T-H Hubert Chan, Rafael Pass, and Elaine Shi. 2018. PaLa: A Simple Partially Synchronous Blockchain *. , 21 pages. <https://eprint.iacr.org/2018/981>
- [24] T-H. Hubert Chan, Rafael Pass, and Elaine Shi. 2018. PiLi: An Extremely Simple Synchronous Blockchain. <https://eprint.iacr.org/2018/980>
- [25] Baibhab Chatterjee, Dong-Hyun Seo, Shramana Chakraborty, Shitij Avlani, Xiaofan Jiang, Heng Zhang, Mustafa Abdallah, Nithin Raghunathan, Charilaos Mousoulis, Ali Shakouri, Saurabh Bagchi, Dimitrios Peroulis, and Shreyas Sen. 2021. Context-Aware Collaborative Intelligence With Spatio-Temporal In-Sensor-Analytics for Efficient Communication in a Large-Area IoT Testbed. *IEEE Internet of Things Journal* 8, 8 (April 2021), 6800–6814. <https://doi.org/10.1109/JIOT.2020.3036087>
- [26] Kim Kwang Raymond Choo, Uttam Ghosh, Deepak Tosh, Reza M. Parizi, and Ali Dehghantanha. 2021. Introduction to the Special Issue on Decentralized Blockchain Applications and Infrastructures for Next Generation Cyber-Physical Systems. *ACM Transactions on Internet Technology* 21, 2 (June 2021), 38e:1–38e:3. <https://doi.org/10.1145/3464768>
- [27] Jeffrey Considine, Matthias Fitzi, Matthew Franklin, Leonid A Levin, Ueli Maurer, and David Metcalf. 2005. Byzantine Agreement Given Partial Broadcast. *Journal of Cryptology* 18, 3 (July 2005), 191–217. <https://doi.org/10.1007/s00145-005-0308-x>
- [28] Ronald Cramer and Victor Shoup. 2000. Signature Schemes Based on the Strong RSA Assumption. *ACM Transactions on Information and System Security* 3, 3 (Aug. 2000), 161–185. <https://doi.org/10.1145/357830.357847>
- [29] Sourav Das, Zhuolun Xiang, and Ling Ren. 2021. Asynchronous Data Dissemination and Its Applications. *Cryptology ePrint Archive* (2021). <https://eprint.iacr.org/2021/777>
- [30] Jérémie Decouchant, David Kozhaya, Vincent Rahli, and Jiangshan Yu. 2022. DAMYSUS: Streamlined BFT Consensus Leveraging Trusted Components. In *Proceedings of the Seventeenth European Conference on Computer Systems (EuroSys '22)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3492321.3519568>
- [31] Dina S. Deif and Yasser Gadallah. 2017. An Ant Colony Optimization Approach for the Deployment of Reliable Wireless Sensor Networks. *IEEE Access* 5 (2017), 10744–10756. <https://doi.org/10.1109/ACCESS.2017.2711484>
- [32] Danny Dolev and H Raymond Strong. 1983. Authenticated Algorithms for Byzantine Agreement. *SIAM J. Comput.* 12, 4 (Nov. 1983), 656–666. <https://doi.org/10.1137/0212045>
- [33] Sisi Duan and Haibin Zhang. 2022. Byzantine Reliable Broadcast with $\$O(nL+kn+n^2 \log n)\$$ Communication. *Cryptology ePrint Archive* (2022). <https://eprint.iacr.org/2022/554>
- [34] Sidi Boubacar ElMamy, Hichem Mrabet, Hassen Gharbi, Abderrazak Jemai, and Damien Trentesaux. 2020. A Survey on the Usage of Blockchain Technology for Cyber-Threats in the Context of Industry 4.0. *Sustainability* 12, 21 (Jan. 2020), 9179. <https://doi.org/10.3390/su12219179>
- [35] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. 2011. Efficient Network Flooding and Time Synchronization with Glossy. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks*. 73–84.
- [36] NIST Fips. 1994. 186 Digital Signature Standard.
- [37] Mattias Fitzi and Ueli Maurer. 2000. From Partial Consistency to Global Broadcast. In *Proceedings of the Annual ACM Symposium on Theory of Computing*. ACM Press, New York, New York, USA, 494–503. <https://doi.org/10.1145/335305.335363>
- [38] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The Bitcoin Backbone Protocol: Analysis and Applications. In *Advances in Cryptology – EUROCRYPT 2015 (Lecture Notes in Computer Science)*, Elisabeth Oswald and Marc Fischlin (Eds.). Springer, Berlin, Heidelberg, 281–310. https://doi.org/10.1007/978-3-662-46803-6_10
- [39] Guy Golan Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. 2019. SBFT: A Scalable and Decentralized Trust Infrastructure. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 568–580. <https://doi.org/10.1109/DSN.2019.00063> arXiv:1804.01626
- [40] Song Guo and Oliver W.W. Yang. 2007. Energy-Aware Multicasting in Wireless Ad Hoc Networks: A Survey and Discussion. *Computer Communications* 30, 9 (June 2007), 2129–2148. <https://doi.org/10.1016/j.comcom.2007.04.006>
- [41] Junxian Huang, Feng Qian, Alexandre Gerber, Z Morley Mao, Subhabrata Sen, and Oliver Spatscheck. 2012. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services - MobiSys '12*. ACM Press, New York, New York, USA, 225. <https://doi.org/10.1145/2307636.2307658>
- [42] Xiaofan Jiang, Heng Zhang, Edgardo Alberto Barsallo Yi, Nithin Raghunathan, Charilaos Mousoulis, Somali Chaterji, Dimitrios Peroulis, Ali Shakouri, and Saurabh Bagchi. 2021. Hybrid Low-Power Wide-Area Mesh Network for IoT Applications. *IEEE Internet of Things Journal* 8, 2 (Jan. 2021), 901–915. <https://doi.org/10.1109/JIOT.2020.3009228>
- [43] Don Johnson, Alfred Menezes, and Scott Vanstone. 2001. The Elliptic Curve Digital Signature Algorithm (ECDSA). *International Journal of Information Security* 1, 1 (Aug. 2001), 36–63. <https://doi.org/10.1007/s102070100002>

- [44] Rüdiger Kapitza, Johannes Behl, Christian Cachin, Tobias Distler, Simon Kuhnle, Seyed Vahid Mohammadi, Wolfgang Schröder-Preikschat, and Klaus Stengel. 2012. CheapBFT: Resource-Efficient Byzantine Fault Tolerance. In *Proceedings of the 7th ACM European Conference on Computer Systems (EuroSys '12)*. Association for Computing Machinery, New York, NY, USA, 295–308. <https://doi.org/10.1145/2168836.2168866>
- [45] Muhammad Samir Khan, Syed Shalan Naqvi, and Nitin H Vaidya. 2019. Exact Byzantine Consensus on Undirected Graphs under Local Broadcast Model. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM, New York, NY, USA, 327–336. <https://doi.org/10.1145/3293611.3331619> arXiv:1903.11677
- [46] Thorsten Kleinjung, Kazumaro Aoki, Jens Franke, Arjen K Lenstra, Emmanuel Thomé, Joppe W Bos, Pierrick Gaudry, Alexander Kruppa, Peter L Montgomery, Dag Arne Osvik, Andrey Timofeev, and Paul Zimmermann. [n. d.]. Factorization of a 768-Bit RSA Modulus. ([n. d.]), 18.
- [47] Chiu Yuen Koo. 2004. Broadcast in Radio Networks Tolerating Byzantine Adversarial Behavior. In *Proceedings of the Annual ACM Symposium on Principles of Distributed Computing*, Vol. 23. ACM Press, New York, New York, USA, 275–282. <https://doi.org/10.1145/1011767.1011807>
- [48] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. 2007. Zyzzyva: Speculative Byzantine Fault Tolerance. In *Proceedings of Twenty-First ACM SIGOPS Symposium on Operating Systems Principles - SOSOP '07*, Vol. 41. ACM Press, New York, New York, USA, 45. <https://doi.org/10.1145/1294261.1294267>
- [49] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems* 4, 3 (July 1982), 382–401. <https://doi.org/10.1145/357172.357176>
- [50] Shancang Li, Shanshan Zhao, Po Yang, Panagiotis Andriotis, Lida Xu, and Qindong Sun. 2019. Distributed Consensus Algorithm for Events Detection in Cyber-Physical Systems. *IEEE Internet of Things Journal* 6, 2 (April 2019), 2299–2308. <https://doi.org/10.1109/JIOT.2019.2906157>
- [51] Johannes Merkle and Manfred Lochter. 2010. *Elliptic Curve Cryptography (ECC) Brainpool Standard Curves and Curve Generation*. Request for Comments RFC 5639. Internet Engineering Task Force. <https://doi.org/10.17487/RFC5639>
- [52] Henrique Moniz. 2020. The Istanbul BFT Consensus Algorithm. <https://doi.org/10.48550/arXiv.2002.03613> arXiv:arXiv:2002.03613
- [53] Roberto Montemanni, L.M. Gambardella, and A.K. Das. 2005. The Minimum Power Broadcast Problem in Wireless Networks: A Simulated Annealing Approach. In *IEEE Wireless Communications and Networking Conference, 2005*, Vol. 4. IEEE, 2057–2062. <https://doi.org/10.1109/WCNC.2005.1424835>
- [54] Pete Mutschler. 2018. Threats to Precision Agriculture. (July 2018). <https://policycommons.net/artifacts/1571846/threats-to-precision-agriculture/2261625/>
- [55] Kartik Nayak, Ling Ren, Elaine Shi, Nitin H Vaidya, and Zhuolun Xiang. 2020. Improved Extension Protocols for Byzantine Broadcast and Agreement. *Leibniz International Proceedings in Informatics, LIPIcs* 179 (Feb. 2020). <https://doi.org/10.4230/LIPIcs.DISC.2020.28> arXiv:2002.11321
- [56] Oscar Novo. 2019. Scalable Access Management in IoT Using Blockchain: A Performance Evaluation. *IEEE Internet of Things Journal* 6, 3 (June 2019), 4694–4701. <https://doi.org/10.1109/JIOT.2018.2879679>
- [57] Christos Profentzas, Magnus Almgren, and Olaf Landsiedel. 2019. IoTLogBlock: Recording Off-line Transactions of Low-Power IoT Devices Using a Blockchain. In *2019 IEEE 44th Conference on Local Computer Networks (LCN)*. 414–421. <https://doi.org/10.1109/LCN44214.2019.8990728>
- [58] Christos Profentzas, Magnus Almgren, and Olaf Landsiedel. 2020. TinyEVM: Off-Chain Smart Contracts on Low-Power IoT Devices. In *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)*. 507–518. <https://doi.org/10.1109/ICDCS47774.2020.00025>
- [59] Sudipta Saha, Olaf Landsiedel, and Mun Choon Chan. 2017. Efficient Many-to-Many Data Sharing Using Synchronous Transmission and TDMA. In *2017 13th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. 19–26. <https://doi.org/10.1109/DCOSS.2017.11>
- [60] Eryk Schiller, Elfat Esati, Sina Rafati Niya, and Burkhard Stiller. 2020. Blockchain on MSP430 with IEEE 802.15.4. In *2020 IEEE 45th Conference on Local Computer Networks (LCN)*. 345–348. <https://doi.org/10.1109/LCN48667.2020.9314805>
- [61] Fred B Schneider. 1990. Implementing Fault-Tolerant Services Using the State Machine Approach: A Tutorial. *Comput. Surveys* 22, 4 (Dec. 1990), 299–319. <https://doi.org/10.1145/98163.98167>
- [62] Yanko Sheiretov, Dave Grundy, Vladimir Zilberstein, Neil Goldfine, and Susan Maley. 2009. MWM-Array Sensors for In Situ Monitoring of High-Temperature Components in Power Plants. *IEEE Sensors Journal* 9, 11 (Nov. 2009), 1527–1536. <https://doi.org/10.1109/JSEN.2009.2019335>
- [63] Nibesh Shrestha, Ittai Abraham, Ling Ren, and Kartik Nayak. 2020. On the Optimality of Optimistic Responsiveness. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. ACM, New York, NY, USA, 839–857. <https://doi.org/10.1145/3372297.3417284>
- [64] M.L. Sichitiu and C. Veerarittiphan. 2003. Simple, Accurate Time Synchronization for Wireless Sensor Networks. In *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003.*, Vol. 2. 1266–1273 vol.2. <https://doi.org/10.1109/WCNC.2003.1200555>
- [65] Atul Singh, Pedro Fonseca, Petr Kuznetsov, Rodrigo Rodrigues, and Petros Maniatis. 2009. Zeno: Eventually Consistent Byzantine-Fault Tolerance.. In *NSDI*, Vol. 9. 169–184.
- [66] Hui Song, Sencun Zhu, and Guohong Cao. 2007. Attack-Resilient Time Synchronization for Wireless Sensor Networks. *Ad Hoc Networks* 5, 1 (Jan. 2007), 112–125. <https://doi.org/10.1016/j.adhoc.2006.05.016>
- [67] Suman Srinivasan, Haniph Latchman, John Shea, Tan Wong, and Janice McNair. 2004. Airborne Traffic Surveillance Systems. In *Proceedings of the ACM 2nd International Workshop on Video Surveillance & Sensor Networks - VSSN '04*. ACM Press, New York, New York, USA, 131. <https://doi.org/10.1145/1026799.1026821>
- [68] Kun Sun, Peng Ning, and Cliff Wang. 2006. TinySeRSync: Secure and Resilient Time Synchronization in Wireless Sensor Networks. In *Proceedings of the 13th ACM Conference on Computer and Communications Security (CCS '06)*. Association for Computing Machinery, New York, NY, USA, 264–277. <https://doi.org/10.1145/1180405.1180439>
- [69] Don Torrieri. 2018. *Principles of Spread-Spectrum Communication Systems*. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-70569-9>
- [70] Giuliana Santos Veronese, Miguel Correia, Alysso Neves Bessani, Lau Cheuk Lung, and Paulo Verissimo. 2013. Efficient Byzantine Fault-Tolerance. *IEEE Trans. Comput.* 62, 1 (Jan. 2013), 16–30. <https://doi.org/10.1109/TC.2011.221>
- [71] Gerald Wagenknecht, Markus Anwender, and Torsten Braun. 2012. SNOMC: An Overlay Multicast Protocol for Wireless Sensor Networks. In *2012 9th Annual Conference on Wireless On-Demand Network Systems and Services (WONS)*. 75–78. <https://doi.org/10.1109/WONS.2012.6152242>
- [72] Chieh-Yih Wan, Andrew T. Campbell, and Lakshman Krishnamurthy. 2002. PSFQ: A Reliable Transport Protocol for Wireless Sensor Networks. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications (WSNA '02)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/570738.570740>
- [73] Shu-Ching Wang, Kuo-Qin Yan, Chin-Ling Ho, and Shun-Sheng Wang. 2014. The Optimal Generalized Byzantine Agreement in Cluster-based Wireless Sensor Networks. *Computer Standards & Interfaces* 36, 5 (Sept. 2014), 821–830. <https://doi.org/10.1016/j.csi.2014.01.005>
- [74] Shun-Sheng Wang, Shu-Ching Wang, and Kuo-Qin Yan. 2014. Reaching Trusted Byzantine Agreement in a Cluster-Based Wireless Sensor Network. *Wireless Personal Communications* 78, 2 (Sept. 2014), 1079–1094. <https://doi.org/10.1007/s11277-014-1802-3>
- [75] Sravya Yandamuri, Ittai Abraham, Kartik Nayak, and Michael K. Reiter. 2023. Communication-Efficient BFT Using Small Trusted Hardware to Tolerate Minority Corruption. In *26th International Conference on Principles of Distributed Systems (OPDIS 2022) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 253)*. Eshcar Hillel, Roberto Palmieri, and Etienne Rivière (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 24:1–24:23. <https://doi.org/10.4230/LIPIcs.OPDIS.2022.24>
- [76] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. 2019. HotStuff: BFT Consensus with Linearity and Responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC '19)*. Association for Computing Machinery, New York, NY, USA, 347–356. <https://doi.org/10.1145/3293611.3331591>
- [77] Feng Yu, Yaodan Hu, Teng Zhang, and Yier Jin. 2020. Special Issue: Resilient Distributed Estimator with Information Consensus for CPS Security. In *2020 IEEE 38th International Conference on Computer Design (ICCD)*. 41–44. <https://doi.org/10.1109/ICCD50377.2020.00023>
- [78] Shanshan Zhao, Shancang Li, and Yufeng Yao. 2019. Blockchain Enabled Industrial Internet of Things Technology. *IEEE Transactions on Computational Social Systems* 6, 6 (Dec. 2019), 1442–1453. <https://doi.org/10.1109/TCSS.2019.2924054>
- [79] Wenbing Zhao, Congfeng Jiang, Honghao Gao, Shunkun Yang, and Xiong Luo. 2021. Blockchain-Enabled Cyber-Physical Systems: A Review. *IEEE Internet of Things Journal* 8, 6 (March 2021), 4023–4034. <https://doi.org/10.1109/JIOT.2020.3014864>
- [80] Marco Zimmerling, Luca Mottola, and Silvia Santini. 2020. Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services. *Comput. Surveys* 53, 6 (Dec. 2020), 121:1–121:39. <https://doi.org/10.1145/3410159>