

# RCC: Resilient Concurrent Consensus for High-Throughput Secure Transaction Processing

Suyash Gupta

Jelle Hellings  
Moka Blox LLC

Mohammad Sadoghi

Exploratory Systems Lab, Department of Computer Science  
University of California, Davis

**Abstract**—Recently, we saw the emergence of consensus-based database systems that promise resilience against failures, strong data provenance, and federated data management. Typically, these fully-replicated systems are operated on top of a primary-backup consensus protocol, which limits the throughput of these systems to the capabilities of a single replica (the primary).

To push throughput beyond this single-replica limit, we propose *concurrent consensus*. In concurrent consensus, replicas independently propose transactions, thereby reducing the influence of any single replica on performance. To put this idea in practice, we propose our RCC paradigm that can turn any primary-backup consensus protocol into a *concurrent consensus protocol* by running many consensus instances concurrently. RCC is designed with performance in mind and requires minimal coordination between instances. Furthermore, RCC also promises increased resilience against failures. We put the design of RCC to the test by implementing it in RESILIENTDB, our high-performance resilient blockchain fabric, and comparing it with state-of-the-art primary-backup consensus protocols. Our experiments show that RCC achieves up to  $2.75\times$  higher throughput than other consensus protocols and can be scaled to 91 replicas.

**Index Terms**—High-throughput resilient transaction processing, concurrent consensus, limits of primary-backup consensus.

## I. INTRODUCTION

Fueled by the emergence of blockchain technology [2], [3], [4], we see a surge in consensus-based data processing frameworks and database systems [4], [5], [6], [7], [8], [9]. This interest can be easily explained: compared to traditional distributed database systems, consensus-based systems can provide *more resilience during failures*, can provide strong support for *data provenance*, and can enable *federated data processing* in a heterogeneous environment with many independent participants. Consequently, consensus-based systems can prevent disruption of service due to software issues or cyberattacks that compromise part of the system, and can aid in improving data quality of data that is managed by many independent parties, potentially reducing the huge societal costs of cyberattacks and bad data.

A brief announcement of this work was presented at the 33rd International Symposium on Distributed Computing (DISC 2019) [1].

This material is based upon work partially supported by the U.S. Department of Energy, Office of Science, Office of Small Business Innovation Research, under Award Number DE-SC0020455.

At the core of consensus-based systems are *consensus protocols* that enable independent participants (e.g., different companies) to manage a single common database by reliably and continuously replicating a unique sequence of transactions among all participants. By design, these consensus protocols are resilient and can deal with participants that have crashed, are unable to participate due to local network, hardware, or software failures, or are compromised and act malicious [10], [11]. As such, consensus protocols can be seen as the fault-resilient counterparts of classical two-phase and three-phase commit protocols [12], [13], [14]. Most practical systems use consensus protocols that follow the classical primary-backup design of PBFT [15] in which a single replica, the *primary*, proposes transactions by broadcasting them to all other replicas, after which all replicas *exchange state* to determine whether the primary correctly proposes the same transaction to all replicas and to deal with failure of the primary. Well-known examples of such protocols are PBFT [15], ZYZZYVA [16], SBFT [17], HOTSTUFF [18], PoE [19], and RBFT [20], and fully-optimized implementations of these protocols are able to process up-to tens-of-thousands transactions per second [21].

### A. The Limitations of Traditional Consensus

Unfortunately, a close look at the design of primary-backup consensus protocols reveals that their design *underutilized available network resources*, which prevents the *maximization of transaction throughput*: the throughput of these protocols is determined mainly by the outgoing bandwidth of the primary. To illustrate this, we consider the maximum throughput by which primaries can replicate transactions. Consider a system with  $\mathbf{n}$  replicas of which  $\mathbf{f}$  are faulty and the remaining  $\mathbf{nf} = \mathbf{n} - \mathbf{f}$  are non-faulty. The maximum throughput  $T_{\max}$  of any such protocol is determined by the outgoing bandwidth  $B$  of the primary, the number of replicas  $\mathbf{n}$ , and the size of transactions  $st$ :  $T_{\max} = B/((\mathbf{n}-1)st)$ . No practical consensus protocol will be able to achieve this throughput, as dealing with crashes and malicious behavior requires substantial state exchange. Protocols such as ZYZZYVA [16] can come close, however, by optimizing for the case in which no faults occur, this at the cost of their ability to deal with faults efficiently.

For PBFT, the minimum amount of state exchange consists of two rounds in which PREPARE and COMMIT messages are exchanged between all replicas (a quadratic amount, see

Example III.1 in Section III). Assuming that these messages have size  $sm$ , the maximum throughput of PBFT is  $T_{\text{PBFT}} = B/((n-1)(st+3sm))$ . To minimize overhead, typical implementations of PBFT group hundreds of transactions together, assuring that  $st \gg sm$  and, hence,  $T_{\text{max}} \approx T_{\text{PBFT}}$ .

The above not only shows a maximum on throughput, but also that primary-backup consensus protocols such as PBFT and ZYZZYVA severely *underutilize resources* of non-primary replicas: when  $st \gg sm$ , the primary sends and receives roughly  $(n-1)st$  bytes, whereas all other replicas only send and receive roughly  $st$  bytes. The obvious solution would be to use several primaries. Unfortunately, recent protocols such as HOTSTUFF [18], SPINNING [22], and PRIME [23] that regularly *switch primaries* all require that a switch from a primary happens after all proposals of that primary are processed. Hence, such primary switching does load balance overall resource usage among the replicas, but does not address the underutilization of resources we observe.

### B. Our Solution: Towards Resilient Concurrent Consensus

The only way to push throughput of consensus-based databases and data processing systems beyond the limit  $T_{\text{max}}$ , is by better utilizing available resources. In this paper, we propose to do so via *concurrent consensus*, in which we use many primaries that *concurrently* propose transactions. We also propose RCC, a paradigm for the realization of concurrent consensus. Our contributions are as follows:

- 1) First, in Section II, we propose *concurrent consensus* and show that concurrent consensus can achieve much higher throughput than primary-backup consensus by effectively utilizing all available system resources.
- 2) Then, in Section III, we propose RCC, a paradigm for turning any primary-backup consensus protocol into a concurrent consensus protocol and that is designed for maximizing throughput in all cases, even during malicious activity.
- 3) Then, in Section IV, we show that RCC can be utilized to make systems *more resilient*, as it can mitigate the effects of order-based attacks and throttling attacks (which are not prevented by traditional consensus protocols), and can provide better load balancing.
- 4) Finally, in Section V, we put the design RCC to the test by implementing it in RESILIENTDB,<sup>1</sup> our high-performance resilient blockchain fabric, and compare RCC with state-of-the-art primary-backup consensus protocols. Our comparison shows that RCC answers the promises of *concurrent consensus*: it achieves up to  $2.75\times$  higher throughput than other consensus protocols, has a peak throughput of  $365\text{ ktxn}/\text{batch}$  and can be easily scaled to 91 replicas.

## II. THE PROMISE OF CONCURRENT CONSENSUS

To deal with the underutilization of resources and the low throughput of primary-backup consensus, we propose

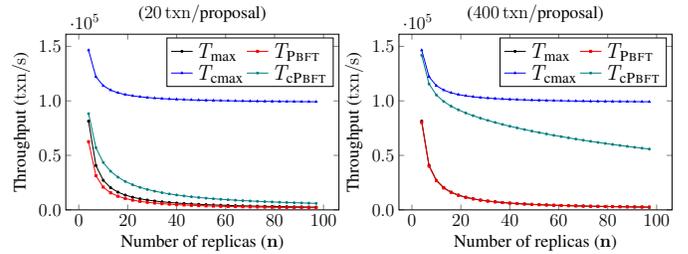


Fig. 1. Maximum throughput of replication in a system with  $B = 1\text{ Gbit/s}$ ,  $n = 3f + 1$ ,  $nf = 2f + 1$ ,  $sm = 1\text{ KiB}$ , and individual transactions are  $512\text{ B}$ . On the *left*, each proposal groups 20 transactions ( $st = 10\text{ KiB}$ ) and on the *right*, each proposal groups 400 transactions ( $st = 2\text{ MiB}$ ).

*concurrent consensus*. In specific, we design for a system that is optimized for high-throughput scenarios in which a plentitude of transactions are available, and we make every replica a *concurrent primary* that is responsible for proposing and replicating some of these transactions. As we have  $nf$  non-faulty replicas, we can expect to always *concurrently* propose at least  $nf$  transactions if sufficient transactions are available. Such concurrent processing has the potential to drastically improve throughput: in each round, each primary will send out one proposal to all other replicas, and receive  $nf - 1$  proposals from other primaries. Hence, the maximum concurrent throughput is  $T_{\text{cmax}} = nfB/((n-1)st + (nf-1)st)$ .

In practice, of course, the primaries also need to participate in state exchange to determine the correct operations of all concurrent primaries. If we use PBFT-style state exchange, we end up with a concurrent throughput of  $T_{\text{cPBFT}} = nfB/((n-1)(st+3sm) + (nf-1)(st+4(n-1)sm))$ . In Figure 1, we have sketched the maximum throughputs  $T_{\text{max}}$ ,  $T_{\text{PBFT}}$ ,  $T_{\text{cmax}}$ , and  $T_{\text{cPBFT}}$ . As one can see, concurrent consensus not only promises greatly improved throughput, but also sharply reduces the costs associated with scaling consensus. We remark, however, that these figures provide best-case *upper-bounds*, as they only focus on bandwidth usage. In practice, replicas are also limited by computational power and available memory buffers that puts limits on the number of transactions they can process in parallel and can execute (see Section V-B).

## III. RCC: RESILIENT CONCURRENT CONSENSUS

The idea behind concurrent consensus, as outlined in the previous section, is straightforward: improve overall throughput by using all available resources via concurrency. Designing and implementing a concurrent consensus system that operates correctly, even during crashes and malicious behavior of some replicas, is challenging, however. In this section, we describe how to design correct *consensus protocols* that deliver on the promises of concurrent consensus. We do so by introducing RCC, a paradigm that can turn any primary-backup consensus protocol into a concurrent consensus protocol. At its basis, RCC makes every replica a primary of a consensus-instance that replicates transactions among all replicas. Furthermore, RCC provides the necessary coordination between these consensus-instances to coordinate execution and deal

<sup>1</sup>RESILIENTDB is open-sourced and available at <https://resilientdb.com>.

with faulty primaries. To assure resilience and maximize throughput, we put the following design goals in RCC:

- D1) RCC provides *consensus* among replicas on the client transactions that are to be executed and the order in which they are executed.
- D2) Clients can interact with RCC to force execution of their transactions and learn the outcome of execution.
- D3) RCC is a design paradigm that can be applied to any primary-backup consensus protocol, turning it into a concurrent consensus protocol.
- D4) In RCC, consensus-instances with non-faulty primaries are *always* able to propose transactions at maximum throughput (with respect to the resources available to any replica), this independent of faulty behavior by any other replica.
- D5) In RCC, dealing with faulty primaries does not interfere with the operations of other consensus-instances.

Combined, design goals D4 and D5 imply that instances with non-faulty primaries can propose transactions *wait-free*: transactions are proposed concurrent to any other activities and does not require any coordination with other instances.

#### A. Background on Primary-Backup Consensus and PBFT

Before we present RCC, we provide the necessary background and notation for primary-backup consensus. Typical primary-backup consensus protocols operate in *views*. Within each view, a primary can propose client transactions, which will then be executed by all non-faulty replicas. To assure that all non-faulty replicas maintain the same state, transactions are required to be *deterministic*: on identical inputs, execution of a transaction must always produce identical outcomes. To deal with faulty behavior by the primary or by any other replicas during a view, three complimentary mechanisms are used:

*Byzantine commit*: The primary uses a *Byzantine commit algorithm* BCA to propose a client transaction  $T$  to all replicas. Next, BCA will perform state exchange to determine whether the primary successfully proposed a transaction. If the primary is non-faulty, then all replicas will receive  $T$  and determine success. If the primary is faulty and more than  $f$  non-faulty replicas do *not* receive a proposal or receive different proposals than the other replicas, then the state exchange step of BCA will detect this failure of the primary.

*Primary replacement*: The replicas use a *view-change algorithm* to replace the primary of the current view  $v$  when this primary is detected to be faulty by non-faulty replicas. This view-change algorithm will collect the state of sufficient replicas in view  $v$  to determine a correct starting state for the next view  $v + 1$  and assign new primary that will propose client transactions in view  $v + 1$ .

*Recovery*: A faulty primary can keep up to  $f$  non-faulty replicas *in the dark* without being detected, as  $f$  faulty replicas can cover for this malicious behavior. Such behavior is not detected and, consequently, does not trigger a view-change. Via a *checkpoint algorithm* the at-most- $f$  non-faulty replicas that are in the dark will learn the proposed client transactions

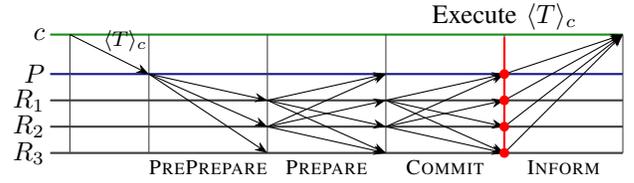


Fig. 2. A schematic representation of the *preprepare-prepare-commit* protocol of PBFT. First, a client  $c$  requests transaction  $T$  and the primary  $P$  proposes  $T$  to all replicas via a PREPREPARE message. Next, replicas commit to  $T$  via a two-phase message exchange (PREPARE and COMMIT messages). Finally, replicas execute the proposal and inform the client.

that are successfully proposed to the remaining at-least- $nf - f > f$  non-fault replicas (that are not in the dark).

**Example III.1.** Next, we illustrate these mechanisms in PBFT. At the core of PBFT is the *preprepare-prepare-commit* Byzantine commit algorithm. This algorithm operates in three phases, which are sketched in Figure 2.

First, the current primary chooses a client request of the form  $\langle T \rangle_c$ , a transaction  $T$  signed by client  $c$ , and proposes this request as the  $\rho$ -th transaction by broadcasting it to all replicas via a PREPREPARE message  $m$ . Next, each non-faulty replica  $R$  *prepares* the first proposed  $\rho$ -th transaction it receives by broadcasting a PREPARE message for  $m$ . If a replica  $R$  receives  $nf$  PREPARE messages for  $m$  from  $nf$  distinct replicas, then it has the guarantee that any group of  $nf$  replicas will contain a *non-faulty replica* that has received  $m$ . Hence,  $R$  has the guarantee that  $m$  can be recovered from any group of  $nf$  replicas, independent of the behavior of the current primary. With this guarantee,  $R$  *commits* to  $m$  by broadcasting a COMMIT message for  $m$ . Finally, if a replica  $R$  receives  $nf$  COMMIT messages for  $m$  from  $nf$  distinct replicas, then it *accepts*  $m$ . In PBFT, accepted proposals are then *executed* and the client is informed of the outcome.

Each replica  $R$  participating in preprepare-prepare-commit uses an internal *timeout value* to detect failure: whenever the primary fails to coordinate a round of preprepare-prepare-commit—which should result in  $R$  accepting some proposal— $R$  will detect *failure* of the primary and halt participation in preprepare-prepare-commit. If  $f + 1$  non-faulty replicas detect such a failure and communication is reliable, then they can cooperate to assure that all non-faulty replicas detect the failure. We call this a *confirmed failure* of preprepare-prepare-commit. In PBFT, confirmed failures trigger a *view-change*. Finally, PBFT employs a majority-vote *checkpoint protocol* that allows replicas that are kept in the dark to learn *accepted* proposals without help of the primary.

#### B. The Design of RCC

We now present RCC in detail. Consider a primary-backup consensus protocol  $P$  that utilizes Byzantine commit algorithm BCA (e.g., PBFT with preprepare-prepare-commit). At the core of applying our RCC paradigm to  $P$  is running  $m$ ,  $1 \leq m \leq n$ , instances of BCA *concurrently*, while providing sufficient coordination between the instances to deal with any

malicious behavior. To do so, RCC makes BCA *concurrent* and uses a checkpoint protocol for per-instance recovery of in-the-dark replicas (see Section III-D). Instead of view-changes, RCC uses a novel wait-free mechanism, that does not involve replacing primaries, to deal with detectable primary failures (see Section III-C). RCC requires the following guarantees on BCA:

**Assumption.** Consider an instance of BCA running in a system with  $n$  replicas,  $n > 3f$ .

- A1) If no failures are detected in round  $\rho$  of BCA (the round is *successful*), then at least  $nf - f$  non-faulty replicas have *accepted* a proposed transaction in round  $\rho$ .
- A2) If a non-faulty replica *accepts* a proposed transaction  $T$  in round  $\rho$  of BCA, then all other non-faulty replicas that accepted a proposed transaction, accepted  $T$ .
- A3) If a non-faulty replica *accepts* a transaction  $T$ , then  $T$  can be recovered from the state of any subset of  $nf - f$  non-faulty replicas.
- A4) If the primary is non-faulty and communication is reliable, then all non-faulty replicas will accept a proposal in round  $\rho$  of BCA.

With minor fine-tuning, these assumptions are met by PBFT, ZYZZYVA, SBFT, HOTSTUFF, and many other primary-backup consensus protocols, meeting design goal D3.

RCC operates in rounds. In each round, RCC replicates  $m$  client transactions (or, as discussed in Section I-A,  $m$  sets of client transactions), one for each instance. We write  $\mathcal{I}_i$  to denote the  $i$ -th instance of BCA. To enforce that each instance is coordinated by a distinct primary, the  $i$ -th replica  $\mathcal{P}_i$  is assigned as the primary coordinating  $\mathcal{I}_i$ . Initially, RCC operates with  $m = n$  instances. In RCC, instances can fail and be *stopped*, e.g., when coordinated by malicious primaries or during periods of unreliable communication. Each round  $\rho$  of RCC operates in three steps:

- 1) *Concurrent BCA.* First, each replica participates in  $m$  instances of BCA, in which each instance is proposing a transaction requested by a client among all replicas.
- 2) *Ordering.* Then, each replica collects all successfully replicated client transactions and puts them in the same—deterministically determined—*order*.
- 3) *Execution.* Finally, each replica *executes* the transactions of round  $\rho$  in order and informs the clients of the outcome of their requested transactions.

Figure 3 sketches a high-level overview of running  $m$  *concurrent instances* of BCA.

To maximize performance, we want every instance to propose distinct transactions, such that every round results in  $m$  distinct transactions. In Section III-E, we delve into the details by which primaries can choose transactions to propose.

To meet design goal D4 and D5, individual BCA instances in RCC can continuously propose and replicate transactions: ordering and execution of the transactions replicated in a round by the  $m$  instances is done *in parallel* to the proposal and replication of transactions for future rounds. Consequently,

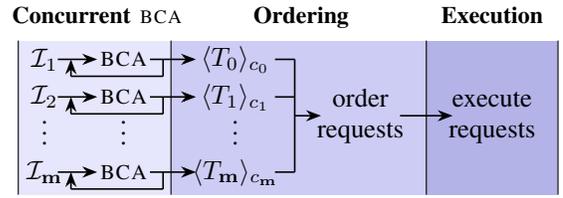


Fig. 3. A high-level overview of RCC running at replica  $R$ . Replica  $R$  participates in  $m$  concurrent instances of BCA (that run independently and continuously output transactions). The instances yield  $m$  transactions, which are executed in a deterministic order.

non-faulty primaries can utilize their entire outgoing network bandwidth for proposing transactions, even if other replicas or primaries are acting malicious.

Let  $\langle T_i \rangle_{c_i}$  be the transaction  $T_i$  requested by  $c_i$  and proposed by  $\mathcal{P}_i$  in round  $\rho$ . After all  $m$  instances complete round  $\rho$ , each replica can collect the set of transactions  $S = \{\langle T_i \rangle_{c_i} \mid 1 \leq i \leq m\}$ . By Assumption A2, all non-faulty replicas will obtain the same set  $S$ . Next, all replicas choose an order on  $S$  and execute all transactions in that order. For now, we assume that the transaction  $\langle T_i \rangle_{c_i}$  is executed as the  $i$ -th transaction of round  $\rho$ . In Section IV, we show that a more advanced ordering-scheme can further improve the resilience of consensus against malicious behavior. As a direct consequence of Assumption A4, we have the following:

**Proposition III.2.** *Consider RCC running in a system with  $n$  replicas,  $n > 3f$ . If all  $m$  instances have non-faulty primaries and communication is reliable, then, in each round, all non-faulty replicas will accept the same set of  $m$  transactions and execute these transactions in the same order.*

As all non-faulty replicas will execute each transaction in  $\langle T_i \rangle_{c_i} \in S$ , there are  $nf$  distinct non-faulty replicas that can inform the client of the outcome of execution. As all non-faulty replicas operate deterministically and execute the transactions in the same order, client  $c_i$  will receive identical outcomes of  $nf > f$  replicas, guaranteeing that this outcome is correct.

In the above, we described the normal-case operations of RCC. As in normal primary-backup protocols, individual instances in RCC can be subject to both *detectable* and *undetectable* failures. Next, we deal with these two types of failures.

### C. Dealing with Detectable Failures

Consensus-based systems typically operate in an environment with asynchronous communication: messages can get lost, arrive with arbitrary delays, and in arbitrary order. Consequently, it is impossible to distinguish between, on the one hand, a primary that is malicious and does not send out proposals and, on the other hand, a primary that does send out proposals that get lost in the network. As such, asynchronous consensus protocols can only provide *progress* in periods of *reliable bounded-delay communication* during which all messages sent by non-faulty replicas will arrive at their destination within some maximum delay [24], [25].

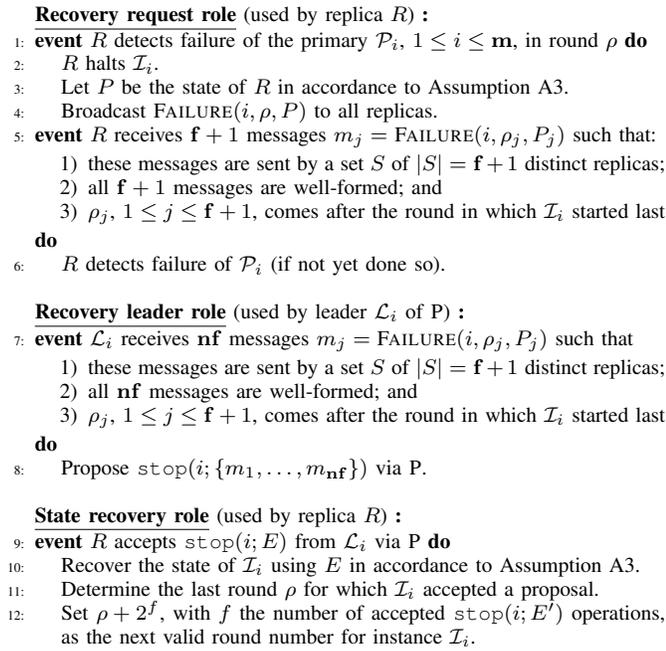


Fig. 4. The recovery algorithm of RCC.

To be able to deal with failures, RCC assumes that *any failure* of non-faulty replicas to receive proposals from a primary  $\mathcal{P}_i$ ,  $1 \leq i \leq \mathbf{m}$ , is due to *failure* of  $\mathcal{P}_i$ , and we design the recovery process such that it can also recover from failures due to unreliable communication. Furthermore, in accordance with the wait-free design goals D4 and D5, the recovery process will be designed so that it does not interfere with other BCA instances or other recovery processes. Now assume that primary  $\mathcal{P}_i$  of  $\mathcal{I}_i$ ,  $1 \leq i \leq \mathbf{m}$ , fails in round  $\rho$ . The recovery process consists of three steps:

- 1) All non-faulty replicas need to detect failure of the  $\mathcal{P}_i$ .
- 2) All non-faulty replicas need to reach agreement on the state of  $\mathcal{I}_i$ : which transactions have been proposed by  $\mathcal{P}_i$  and have been accepted in the rounds up-to- $\rho$ .
- 3) To deal with unreliable communication, all non-faulty replicas need to determine the round in which  $\mathcal{P}_i$  is allowed to resume its operations.

To reach agreement on the state of  $\mathcal{I}_i$ , we rely on a separate instance of the consensus protocol  $P$  that is only used to coordinate agreement on the state of  $\mathcal{I}_i$  during failure. This coordinating consensus protocol  $P$  replicates  $\text{stop}(i; E)$  operations, in which  $E$  is a set of  $\mathbf{nf}$  FAILURE messages sent by  $\mathbf{nf}$  distinct replicas from which all accepted proposals in instance  $\mathcal{I}_i$  can be derived. We notice that  $P$  is—itself—an instance of a primary-backup protocol that is coordinated by some primary  $\mathcal{L}_i$  (based on the current view in which the instance of  $P$  operates), and we use the standard machinery of  $P$  to deal with failures of that leader (see Section III-A). Next, we shall describe how the recovery process is initiated. The details of this protocol can be found in Figure 4.

When a replica  $R$  detects failure of instance  $\mathcal{I}_i$ ,  $0 \leq i < \mathbf{m}$ , in round  $\rho$ , it broadcasts a message FAILURE( $i, \rho, P$ ), in

which  $P$  is the state of  $R$  in accordance to Assumption A3 (Line 1 of Figure 4). To deal with unreliable communication,  $R$  will continuously broadcast this FAILURE message with an exponentially-growing delay until it learns on how to proceed with  $\mathcal{I}_i$ . To reduce communication in the normal-case operations of  $P$ , one can send the full message FAILURE( $i, \rho, P$ ) to only  $\mathcal{L}_i$ , while sending FAILURE( $i, \rho$ ) to all other replicas.

If a replica receives  $\mathbf{f} + 1$  FAILURE messages from distinct replicas for a certain instance  $\mathcal{I}_i$ , then it received at least one such message from a non-faulty replica. Hence, it can detect failure of  $\mathcal{I}_i$  (Line 5 of Figure 4). Finally, if a replica  $R$  receives  $\mathbf{nf}$  FAILURE messages from distinct replicas for a certain instance  $\mathcal{I}_i$ , then we say there is a *confirmed failure*, as  $R$  has the guarantee that eventually—within at most two message delays—also the primary  $\mathcal{L}_i$  of  $P$  will receive  $\mathbf{nf}$  FAILURE messages (if communication is reliable). Hence, at this point,  $R$  sets a timer based on some internal timeout value (that estimates the message delay) and waits on the leader  $\mathcal{L}_i$  to propose a valid  $\text{stop}$ -operation or for the timer to run out. In the latter case, replica  $R$  detects failure of the leader  $\mathcal{L}_i$  and follows the steps of a view-change in  $P$  to (try to) replace  $\mathcal{L}_i$ . When the leader  $\mathcal{L}_i$  receives  $\mathbf{nf}$  FAILURE messages, it can and must construct a valid  $\text{stop}$ -operation and reach consensus on this operation (Line 7 of Figure 4). After reaching consensus, each replica can recover to a common state of  $\mathcal{I}_i$ :

**Theorem III.3.** *Consider RCC running in a system with  $\mathbf{n}$  replicas. If  $\mathbf{n} > 3\mathbf{f}$ , an instance  $\mathcal{I}_i$ ,  $0 \leq i < \mathbf{m}$ , has a confirmed failure, and the last proposal of  $\mathcal{P}_i$  accepted by a non-faulty replica was in round  $\rho$ , then—whenever communication becomes reliable—the recovery protocol of Figure 4 will assure that all non-faulty replicas will recover the same state, which will include all proposals accepted by non-faulty replicas before-or-at round  $\rho$ .*

*Proof.* If communication is reliable and instance  $\mathcal{I}_i$  has a confirmed failure, then all non-faulty replicas will detect this failure and send FAILURE messages (Line 1 of Figure 4). Hence, all replicas are guaranteed to receive at least  $\mathbf{nf}$  FAILURE messages, and any replica will be able to construct a well-formed operation  $\text{stop}(i; E)$ . Hence,  $P$  will eventually be forced to reach consensus on  $\text{stop}(i; E)$ . Consequently, all non-faulty replicas will conclude on the same state for instance  $\mathcal{I}_i$ . Now consider a transaction  $T$  accepted by non-faulty replica  $Q$  in instance  $\mathcal{I}_i$ . Due to Assumption A3,  $Q$  will only accept  $T$  if  $T$  can be recovered from the state of any set of  $\mathbf{nf} - \mathbf{f}$  non-faulty replicas. As  $|E| = \mathbf{nf}$  (Line 7 of Figure 4), the set  $E$  contains the state of  $\mathbf{nf} - \mathbf{f}$  non-faulty replicas. Hence,  $T$  must be recoverable from  $E$ .  $\square$

We notice that the recovery algorithm of RCC, as outlined in Figure 4, only affects the capabilities of the BCA instance that is stopped. All other BCA instances can concurrently propose transactions for current and for future rounds. Hence, the recovery algorithm adheres to the wait-free design goals D4 and D5. Furthermore, we reiterate that we have separate instance of the coordinating consensus protocol for

each instance  $\mathcal{I}_i$ ,  $1 \leq i \leq m$ . Hence, recovery of several instances can happen concurrently, which minimizes the time it takes to recover from several simultaneous primary failures and, consequently, minimizes the delay before a round can be executed during primary failures.

Confirmed failures not only happen due to malicious behavior. Instances can also fail due to periods of unreliable communication. To deal with this, we eventually restart any stopped instances. To prevent instances coordinated by malicious replicas to continuously cause recovery of their instances, every failure will incur an exponentially growing restart penalty (Line 12 of Figure 4). The exact round in which an instance can resume operations can be determined deterministically from the accepted history of `stop`-requests. When all instances have round failures due to unreliable communication (which can be detected from the history of `stop`-requests), any instance is allowed to resume operations in the earliest available round (after which all other instances are also required to resume operations).

#### D. Dealing with Undetectable Failures

As stated in Assumption A1, a malicious primary  $\mathcal{P}_i$  of a BCA instance  $\mathcal{I}_i$  is able to keep up to  $f$  non-faulty replicas in the dark without being detected. In normal primary-backup protocols, this is not a huge issue: at least  $nf - f > f$  non-faulty replicas still accept transactions, and these replicas can execute and reliably inform the client of the outcome of execution. This is not the case in RCC, however:

**Example III.4.** Consider a system with  $n = 3f + 1 = 7$  replicas. Assume that primaries  $\mathcal{P}_1$  and  $\mathcal{P}_2$  are malicious, while all other primaries are non-faulty. We partition the non-faulty replicas into three sets  $A_1$ ,  $A_2$ , and  $B$  with  $|A_1| = |A_2| = f$  and  $|B| = 1$ . In round  $\rho$ , the malicious primary  $\mathcal{P}_i$ ,  $i \in \{1, 2\}$ , proposes transaction  $\langle T_i \rangle_{c_i}$  to only the non-faulty replicas in  $A_i \cup B$ . This situation is sketched in Figure 5. After all concurrent instances of BCA finish round  $\rho$ , we see that the replicas in  $A_1$  have accepted  $\langle T_1 \rangle_{c_1}$ , the replicas in  $A_2$  have accepted  $\langle T_2 \rangle_{c_2}$ , and only the replica in  $B$  has accepted both  $\langle T_1 \rangle_{c_1}$  and  $\langle T_2 \rangle_{c_2}$ . Hence, only the single replica in  $B$  can proceed with execution of round  $\rho$ . Notice that, due to Assumption A1, we consider all instances as finished successfully. If  $n \geq 10$  and  $f \geq 3$ , this example attack can be generalized such that also the replica in  $B$  is missing at least a single client transaction.

To deal with *in-the-dark* attacks of Example III.4, we can run a standard checkpoint algorithm for each BCA instance: if the system does not reach confirmed failure of  $\mathcal{P}_i$  in round  $\rho$ ,  $1 \leq i \leq m$ , then, by Assumption A1 and A2, at-least- $nf - f$  non-faulty replicas have accepted the same transaction  $T$  in round  $\rho$  of  $\mathcal{I}_i$ . Hence, by Assumption A3, a standard checkpoint algorithm (e.g., the one of PBFT or one based on delayed replication [26]) that exchanges the state of these at-least- $nf - f$  non-faulty replicas among all other replicas is sufficient to assure that all non-faulty replicas eventually accept  $T$ . We notice that these checkpoint algorithms can be

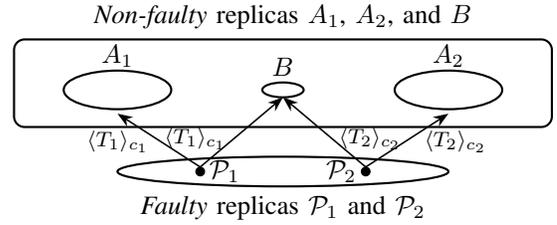


Fig. 5. An attack possible when parallelizing BCA: malicious primaries can prevent non-faulty replicas from learning all client requests in a round, thereby preventing timely round execution. The faulty primary  $\mathcal{P}_i$ ,  $i \in \{1, 2\}$ , does so by only letting non-faulty replicas  $A_i \cup B$  participate in instance  $\mathcal{I}_i$ .

run concurrently with the operations of BCA instances, thereby adhering to our wait-free design goals D4 and D5.

To reduce the cost of checkpoints, typical consensus systems only perform checkpoints after every  $x$ -th round for some system-defined constant  $x$ . Due to in-the-dark attacks, applying such a strategy to RCC means choosing between execution latency and throughput. Consequently, in RCC we do checkpoints on a dynamic *per-need basis*: when replica  $R$  receives  $nf - f$  claims of failure of primaries (via the FAILURE messages of the recovery protocol) in round  $\rho$  and  $R$  itself finished round  $\rho$  for all its instances, then it will participate in any attempt for a checkpoint for round  $\rho$ . Hence, if an in-the-dark attack affects more than  $f$  distinct non-faulty replicas in round  $\rho$ , then a successful checkpoint will be made and all non-faulty replicas recover from the attack, accept all transactions in round  $\rho$ , and execute all these transactions.

Using Theorem III.3 to deal with detectable failures and using checkpoint protocols to deal with replicas in-the-dark, we conclude that RCC adheres to design goal D1:

**Theorem III.5.** Consider RCC running in a system with  $n$  replicas. If  $n > 3f$ , then RCC provides consensus in periods in which communication is reliable.

#### E. Client Interactions with RCC

To maximize performance, it is important that every instance proposes distinct client transactions, as proposing the same client transaction several times would reduce throughput. We have designed RCC with faulty clients in mind, hence, we do not expect cooperation of clients to assure that they send their transactions to only a single primary.

To be able to do so, the design of RCC is optimized for the case in which there are always many more concurrent clients than replicas in the system. In this setting, we assign every client  $c$  to a single primary  $\mathcal{P}_i$ ,  $1 \leq i \leq m = n$ , such that only instance  $\mathcal{I}_i$  can propose client requests of  $c$ . For this design to work in *all cases*, we need to solve two issues, however: we need to deal with situations in which primaries do not receive client requests (e.g., during downtime periods in which only few transactions are requested), and we need to deal with faulty primaries that refuse to propose requests of some clients.

First, if there are less concurrent clients than replicas in the system, e.g., when demand for services is low, then RCC still

needs to process client transactions correctly, but it can do so without optimally utilizing resources available, as this would not impact throughput in this case due to the low demands. If a primary  $\mathcal{P}_i$ ,  $1 \leq i \leq \mathbf{m}$ , does not have transactions to propose in any round  $\rho$  and  $\mathcal{P}_i$  detects that other BCA instances are proposing for round  $\rho$  (e.g., as it receives proposals), then  $\mathcal{P}_i$  proposes a small no-op-request instead.

Second, to deal with a primary  $\mathcal{P}_i$ ,  $1 \leq i \leq \mathbf{m}$ , that refuses to propose requests of some clients, we take a two-step approach. First, we incentivize malicious primaries to *not refuse* services, as otherwise they will be detected faulty and lose the ability to propose transactions altogether. To detect failure of  $\mathcal{P}_i$ , RCC uses standard techniques to enable a client  $c$  to *force* execution of a transaction  $T$ . First,  $c$  broadcasts  $\langle T \rangle_c$  to all replicas. Each non-faulty replica  $R$  will then forward  $\langle T \rangle_c$  to the appropriate primary  $\mathcal{P}_i$ ,  $1 \leq i \leq \mathbf{m}$ . Next, if the primary  $\mathcal{P}_i$  does not propose any transaction requested by  $c$  within a reasonable amount of time, then  $R$  detects failure of  $\mathcal{P}_i$ . Hence, refusal of  $\mathcal{P}_i$  to propose  $\langle T \rangle_c$  will lead to primary failure, incentivizing malicious primaries to provide service.

Finally, we need to deal with primaries that are unwilling or incapable of proposing requests of  $c$ , e.g., when the primary crashes. To do so,  $c$  can request to be reassigned to another instance  $\mathcal{I}_j$ ,  $1 \leq j \leq \mathbf{m}$ , by broadcasting a request  $m := \text{SWITCHINSTANCE}(c, j)$  to all replicas. Reassignment is handled by the coordinating consensus protocol  $\mathcal{P}$  for  $\mathcal{I}_i$ , that will reach consensus on  $m$ . Malicious clients can try to use reassignment to propose transactions in several instances at the same time. To deal with this, we assume that no instance is more than  $\sigma$  rounds behind any other instance (see Section IV). Now, consider the moment at which replica  $R$  accepts  $m$  and let  $\rho(m, R)$  be the maximum round in which any request has been proposed by any instance in which  $R$  participates. The primary  $\mathcal{P}_i$  will stop proposing transactions of  $c$  immediately. Any non-faulty replica  $R$  will stop accepting transactions of  $c$  by  $\mathcal{I}_i$  after round  $\rho(m, R) + \sigma$  and will start accepting transactions of  $c$  by  $\mathcal{I}_j$  after round  $\rho(m, R) + 2\sigma$ . Finally,  $\mathcal{P}_j$  will start proposing transactions of  $c$  in round  $\rho(m, \mathcal{P}_j) + 3\sigma$ .

#### IV. RCC: IMPROVING RESILIENCE OF CONSENSUS

Traditional primary-backup consensus protocols rely heavily on the operations of their primary. Although these protocols are designed to deal with primaries that *completely fail* proposing client transactions, they are not designed to deal with many other types of malicious behavior.

**Example IV.1.** Consider a financial service running on a traditional PBFT consensus-based system. In this setting, a malicious primary can affect operations in two malicious ways:

- 1) *Ordering attack.* The primary sets the order in which transactions are processed and, hence, can choose an ordering that best fits its own interests. To illustrate this, we consider client transactions of the form:

$$\text{transfer}(A, B, n, m) := \text{if amount}(A) > n \text{ then} \\ \text{withdraw}(A, m); \text{ deposit}(B, m).$$

	Original Balance	First $T_1$ , then $T_2$		First $T_2$ , then $T_1$	
		$T_1$	$T_2$	$T_2$	$T_1$
Alice	800	600	600	800	600
Bob	300	500	200	300	500
Eve	100	100	400	100	100

Fig. 6. Illustration of the influence of execution order on the outcome: switching around requests affects the transfer of  $T_2$ .

Let  $T_1 = \text{transfer}(\text{Alice}, \text{Bob}, 500, 200)$  and  $T_2 = \text{transfer}(\text{Bob}, \text{Eve}, 400, 300)$ . Before processing these transactions, the balance for Alice is 800, for Bob 300, and for Eve 100. In Figure 6, we summarize the results of either first executing  $T_1$  or first executing  $T_2$ . As is clear from the figure, execution of  $T_1$  influences the outcome of execution of  $T_2$ . As primaries choose the ordering of transactions, a malicious primary can choose an ordering whose outcome benefits its own interests, e.g., formulate targeted attacks to affect the execution of the transaction of some clients.

- 2) *Throttling attack.* The primary sets the pace at which the system processes transactions. We recall that individual replicas rely on *time-outs* to detect malicious behavior of the primary. This approach will fail to detect or deal with primaries that *throttle throughput* by proposing transactions as slow as possible, while preventing failure detection due to time-outs.

Besides malicious primaries, also other malicious entities can take advantage of a primary-backup consensus protocol:

- 3) *Targeted attack.* As the throughput of a primary-backup system is entirely determined by the primary, attackers can send arbitrary messages to the primary. Even if the primary recognizes that these messages are irrelevant for its operations, it has spend resources (network bandwidth, computational power, and memory) to do so, thereby reducing throughput. Notice that—in the worst case—this can even lead to failure of a non-faulty primary to propose transactions in a timely manner.

Where traditional consensus-based systems fail to deal with these attacks, the concurrent design of RCC can be used to mitigate these attacks.

First, we look at *ordering attacks*. To mitigate this type of attack, we propose a method to deterministically select a different permutation of the order of execution in every round in such a way that this ordering is practically impossible to predict or influence by faulty replicas. Note that for any sequence  $S$  of  $k = |S|$  values, there exist  $k!$  distinct permutations. We write  $P(S)$  to denote these permutations of  $S$ . To deterministically select one of these permutations, we construct a function that maps an integer  $h \in \{0, \dots, k!-1\}$  to a unique permutation in  $P(S)$ . Then we discuss how replicas will uniformly pick  $h$ . As  $|P(S)| = k!$ , we can construct the

following bijection  $f_S : \{0, \dots, k! - 1\} \rightarrow P(S)$

$$f_S(i) = \begin{cases} S & \text{if } |S| = 1; \\ f_{S \setminus S[q]}(r) \oplus S[q] & \text{if } |S| > 1, \end{cases}$$

in which  $q = i \operatorname{div} (|S| - 1)!$  is the quotient and  $r = i \operatorname{mod} (|S| - 1)!$  is the remainder of integer division by  $(|S| - 1)!$ . Using induction on the size of  $S$ , we can prove:

**Lemma IV.2.**  $f_S$  is a bijection from  $\{0, \dots, |S|! - 1\}$  to all possible permutations of  $S$ .

Let  $S$  be the sequence of all transactions accepted in round  $\rho$ , ordered on increasing instance. The replicas uniformly pick  $h = \operatorname{digest}(S) \operatorname{mod} (k! - 1)$ , in which  $\operatorname{digest}(S)$  is a *strong cryptographic hash function* that maps an arbitrary value  $v$  to a numeric digest value in a bounded range such that it is practically impossible to find another value  $S'$ ,  $S \neq S'$ , with  $\operatorname{digest}(S) = \operatorname{digest}(S')$ . When at least one primary is non-malicious ( $\mathbf{m} > \mathbf{f}$ ), the final value  $h$  is only known after completion of round  $\rho$  and it is practically impossible to predictably influence this value. After selecting  $h$ , all replicas execute the transactions in  $S$  in the order given by  $f_S(h)$ .

To deal with primaries that throttle their instances, non-faulty replicas will detect failure of those instances that lag behind other instances. In specific, if an instance  $\mathcal{I}_i$ ,  $1 \leq i \leq \mathbf{m}$ , is  $\sigma$  rounds behind any other instances (for some system-dependent constant  $\sigma$ ), then  $R$  detects failure of  $\mathcal{P}_i$ .

Finally, we notice that concurrent consensus and RCC—by design—provides *load balancing* with respect to the tasks of the primary, this by spreading the total workload of the system over many primaries. As such, RCC not only improves performance when bounded by the primary bandwidth, but also when performance is bounded by computational power (e.g., due to costly cryptographic primitives), or by message delays. Furthermore, this *load balancing* reduces the load on any single primary to propose and process a given amount of transactions, dampening the effects of any targeted attacks against the resources of a single primary.

## V. EVALUATION OF THE PERFORMANCE OF RCC

In the previous sections, we proposed concurrent consensus and presented the design of RCC, our concurrent consensus paradigm. To show that concurrent consensus not only provides benefits in theory, we study the performance of RCC and the effects of concurrent consensus in a practical setting. To do so, we measure the performance of RCC in RESILIENTDB—our high-performance resilient blockchain fabric—and compare RCC with the well-known primary-backup consensus protocols PBFT, ZYZZYVA, SBFT, and HOTSTUFF. With this study, we aim to answer the following questions:

- Q1) What is the performance of RCC: does RCC deliver on the promises of concurrent consensus and provide more throughput than any primary-backup consensus protocol can provide?
- Q2) What is the scalability of RCC: does RCC deliver on the promises of concurrent consensus and provide better scalability than primary-backup consensus protocols?

- Q3) Does RCC provide sufficient load balancing of primary tasks to improve performance of consensus by offsetting any high costs incurred by the primary?
- Q4) How does RCC fare under failures?
- Q5) What is the impact of batching client transactions on the performance of RCC?

First, in Section V-A, we describe the experimental setup. Then, in Section V-B, we provide a high-level overview of RESILIENTDB and of its general performance characteristics. Next, in Section V-C, we provide details on the consensus protocols we use in this evaluation. Then, in Section V-D, we present the experiments we performed and the measurements obtained. Finally, in Section V-E, we interpret these measurements and answer the above research questions.

### A. Experimental Setup

To be able to study the practical performance of RCC and other consensus protocols, we choose to study these protocols in a full resilient database system. To do so, we implemented RCC in RESILIENTDB. To generate a workload for the protocols, we used the *Yahoo Cloud Serving Benchmark* [27] provided by the Blockbench macro benchmarks [28]. In the generated workload, each client transaction queries a YCSB table with half a million active records and 90% of the transactions write and modify records. Prior to the experiments, each replica is initialized with an identical copy of the YCSB table. We perform all experiments in the Google Cloud. In specific, each replica is deployed on a c2-machine with a 16-core Intel Xeon Cascade Lake CPU running at 3.8 GHz and with 32 GB memory. We use up to 320 k clients, deployed on 16 machines.

### B. The RESILIENTDB Blockchain Fabric

The RESILIENTDB fabric incorporates secure permissioned blockchain technologies to provide resilient data processing. A detailed description of how RESILIENTDB achieves high-throughput consensus in a practical settings can be found in Gupta et al. [21], [29], [30], [31], [32]. The architecture of RESILIENTDB is optimized for maximizing throughput via *multi-threading* and *pipelining*. To further maximize throughput and minimize the overhead of any consensus protocol, RESILIENTDB has built-in support for *batching* of client transactions.

We typically group 100 txn/batch. In this case, the size of a proposal is 5400 B and of a client reply (for 100 transactions) is 1748 B. The other messages exchanged between replicas during the Byzantine commit algorithm have a size of 250 B. RESILIENTDB supports *out-of-order processing* of transactions in which primaries can propose future transactions before current transactions are executed. This allows RESILIENTDB to maximize throughput of any primary-backup protocol that supports out-of-order processing (e.g., PBFT, ZYZZYVA, and SBFT) by maximizing bandwidth utilization at the primary.

In RESILIENTDB, each replica maintains a *blockchain ledger* (a journal) that holds an ordered copy of all executed transactions. The ledger not only stores all transactions, but

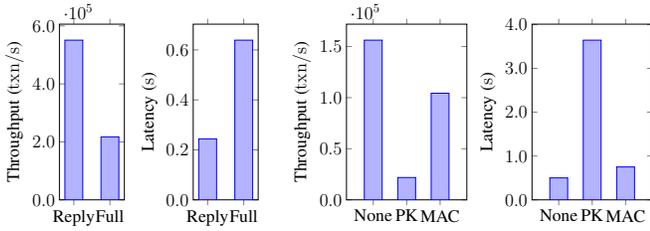


Fig. 7. Characteristics of RESILIENTDB deployed on the Google Cloud. *Left*, the maximum performance of a single replica that receives clients transactions, optionally executes them (Full), and sends replies. *Right*, the performance of PBFT with  $n = 16$  replicas that uses no cryptography (None), uses ED25519 public-key cryptography (PK), or CMAC-AES message authentication codes (MAC) to authenticate messages.

also proofs of their acceptance by a consensus protocols. As these proofs are built using strong cryptographic primitives, the ledger is *immutable* and, hence, can be used to provide *strong data provenance*.

In our experiments replicas not only perform consensus, but also communicate with clients and execute transactions. In this practical setting, performance is not fully determined by bandwidth usage due to consensus (as outlined in Section I-A), but also by the cost of *communicating with clients*, of sequential *execution* of all transactions, of *cryptography*, and of other steps involved in processing messages and transactions, and by the available memory limitations. To illustrate this, we have measured the effects of *client communication*, *execution*, and *cryptography* on our deployment of RESILIENTDB.

In Figure 7, *left*, we present the maximum performance of a single replica that receives clients transactions, optionally executes them (*Full*), and sends replies (without any consensus steps). In this figure, we count the total number of client transactions that are completed during the experiment. As one can see, the system can receive and respond to up-to-551 ktxn/sec, but can only execute up-to-217 ktxn/sec.

In Figure 7, *right*, we present the maximum performance of PBFT running on  $n = 16$  replicas as a function of the cryptographic primitives used to provide *authenticated communication*. In specific, PBFT can either use digital signatures or message authentication codes. For this comparison, we compare PBFT using: (1) a baseline that does not use any message authentication (*None*); (2) ED25519 digital signatures for all messages (*DS*); and (3) CMAC+AES message authentication codes for all messages exchanged between messages and ED25519 digital signatures for client transactions. As can be seen from the results, the costs associated with digital signatures are huge, as their usage reduces performance by 86%, whereas message authentication codes only reduce performance by 33%.

### C. The Consensus Protocols

We evaluate the performance of RCC by comparing it with a representative sample of efficient practical primary-backup consensus protocols:

PBFT [15]: We use a heavily optimized out-of-order implementation that uses message authentication codes.

RCC: Our RCC implementation follows the design outlined in this paper. We have chosen to turn PBFT into a concurrent consensus protocol. We test with three variants:  $RCC_n$  runs  $n$  concurrent instances,  $RCC_{f+1}$  runs  $f + 1$  concurrent instances (the minimum to provide the benefits outlined in Section IV), and  $RCC_3$  runs 3 concurrent instances.

ZYZZYVA [16]: As described in Section I-A, ZYZZYVA has a optimal-case path due to which the performance of ZYZZYVA provides an upper-bound for any primary-backup protocol (when no failures occur). Unfortunately, the failure-handling of ZYZZYVA is costly, making ZYZZYVA unable to deal with any failures efficiently.

SBFT [17]: This protocol uses threshold signatures to minimize communication during the state exchange that is part of its Byzantine commit algorithm. Threshold signatures do not reduce the communication costs for the primary to propose client transactions, which have a major influence on performance in practice (See Section I-A), but can potentially greatly reduce all other communication costs.

HOTSTUFF [18]: As SBFT, HOTSTUFF uses threshold signatures to minimize communication. The state-exchange of HOTSTUFF has an extra phase compared to PBFT. This additional phase simplifies changing views in HOTSTUFF, and enables HOTSTUFF to regularly switch primaries (which limits the influence of any faulty replicas). Due to this design, HOTSTUFF does not support out-of-order processing (see Section I-A). As a consequence, HOTSTUFF is more affected by message delays than by bandwidth. In our implementation, we have used the efficient single-phase event-based variant of HOTSTUFF.

### D. The Experiments

To be able to answer Question Q1–Q5, we perform four experiments in which we measure the performance of RCC. In each experiment, we measure the *throughput* as the number of transactions that are executed per second, and we measure the *latency* as the time from when a client sends a transaction to the time where that client receives a response. We run each experiment for 180s: the first 60s are warm-up, and measurement results are collected over the next 120s. We average our results over three runs. The results of all four experiments can be found in Figure 8.

In the first experiment, we measure the *best-case performance* of the consensus protocols as a function of the number of replicas when all replicas are non-faulty. We vary the number of replicas between  $n = 4$  and  $n = 91$  and we use a batch size of 100 txn/batch. The results can be found in Figure 8, (a) and (b).

In the second experiment, we measure the performance of the consensus protocols as a function of the number of replicas during failure of a single replica. Again, we vary the number of replicas between  $n = 4$  and  $n = 91$  and we use a batch size of 100 txn/batch. The results can be found in Figure 8, (c) and (d).

In the third experiment, we measure the performance of the consensus protocols as a function of the number of replicas

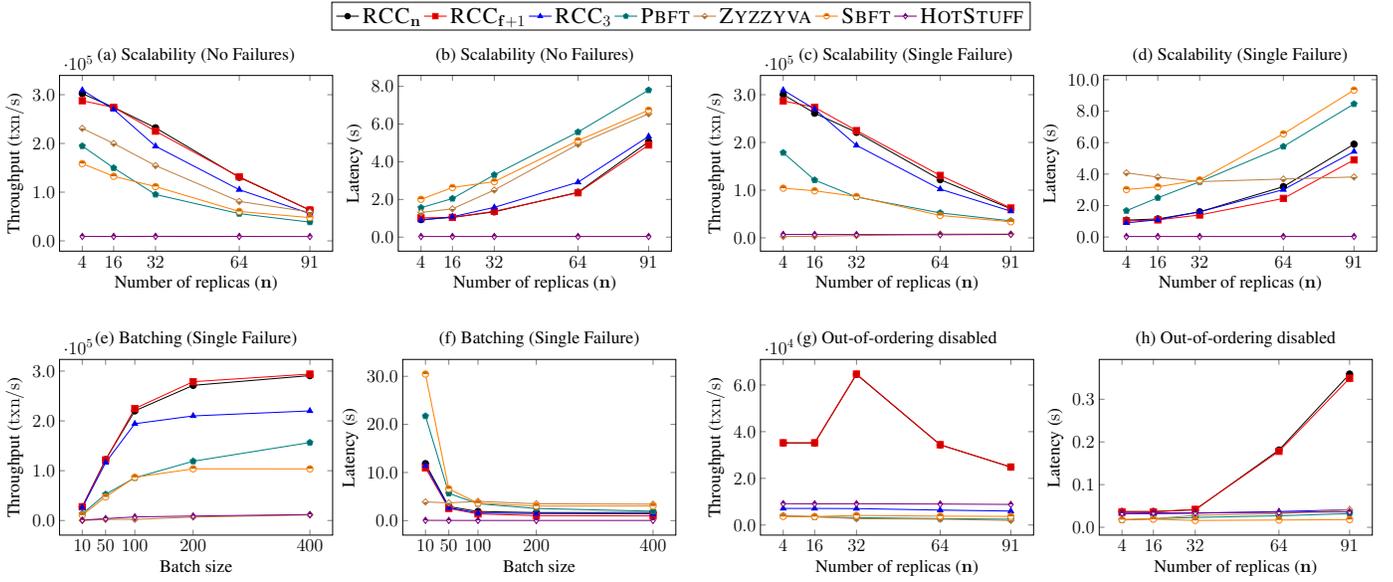


Fig. 8. Evaluating system throughput and average latency incurred by RCC and other consensus protocols.

during failure of a single replica while varying the batch size between 10 txn/batch and 400 txn/batch. We use  $n = 32$  replicas. The results can be found in Figure 8, (e) and (f).

In the fourth and final experiment, we measure the performance of the consensus protocols when outgoing primary bandwidth is not the limiting factor. We do so by disabling *out-of-order processing* in all protocols that support out-of-order processing. This makes the performance of these protocols inherently bounded by the message delay and not by network bandwidth. We study this case by varying the number of replicas between  $n = 4$  and  $n = 91$  and we use a batch size of 100 txn/batch. The results can be found in Figure 8, (g) and (h).

### E. Discussion

From the experiments, a few obvious patterns emerge. First, we see that increasing the batch size ((e) and (f)) increases performance of all consensus protocols (Q5). This is in line with what one can expect (See Section I-A and Section II). As the gains beyond 100 txn/batch are small, we have chosen to use 100 txn/batch in all other experiments.

Second, we see that the three versions of RCC outperform all other protocols, and the performance of RCC with or without failures is comparable ((a)–(d)). Furthermore, we see that adding concurrency by adding more instances improves performance, as  $RCC_3$  is outperformed by the other RCC versions. On small deployments with  $n = 4, \dots, 16$  replicas, the strength of RCC is most evident, as our RCC implementations approach the maximum rate at which RESILIENTDB can execute transactions (see Section V-B).

Third, we see that RCC easily outperforms ZYZZYVA, even in the best-case scenario of no failures ((a) and (b)). We also see that ZYZZYVA is—indeed—the fastest primary-backup consensus protocol when no failures happen. This underlines the ability of RCC, and of concurrent consensus in general, to reach throughputs no primary-backup consensus protocol can

reach. We also notice that ZYZZYVA fails to deal with failures ((c) and (d)), in which case its performance plummets, a case that the other protocols have no issues dealing with.

Finally, due to the lack of out-of-order processing capabilities in HOTSTUFF, HOTSTUFF is uncompetitive to out-of-order protocols. When we disable out-of-order processing for all other protocols ((g) and (h)), the strength of the simple design of HOTSTUFF shows: its event-based single-phase design outperforms all other primary-backup consensus protocols. Due to the concurrent design of RCC, a non-out-of-order-RCC is still able to greatly outperform HOTSTUFF, however, as the non-out-of-order variants of RCC balance the entire workload over many primaries. Furthermore, as the throughput is not bound by any replica resources in this case (and only by network delays), the non-out-of-order variants  $RCC_{f+1}$  and  $RCC_n$  benefit from increasing the number of replicas, as this also increases the amount of concurrent processing (due to increasing the number of instances).

*Summary:* RCC implementations achieve up to  $2.77\times$ ,  $1.53\times$ ,  $38\times$ , and  $82\times$  higher throughput than SBFT, PBFT, HOTSTUFF, and ZYZZYVA in single failure experiments. RCC implementations achieve up to  $2\times$ ,  $1.83\times$ ,  $33\times$ , and  $1.45\times$  higher throughput than SBFT, PBFT, HOTSTUFF, and ZYZZYVA in no failure experiments, respectively.

Based on these observations, we conclude that RCC delivers on the promises of concurrent consensus. RCC provides more throughput than any primary-backup consensus protocol can provide (Q1). Moreover, RCC provides great scalability if throughput is only bounded by the primaries: as the non-out-of-order results show, the load-balancing capabilities of RCC can even offset inefficiencies in other parts of the consensus protocol (Q2, Q3). Finally, we conclude that RCC can efficiently deal with failures (Q4). Hence, RCC meets the design goals D1–D5 that we set out in Section III.

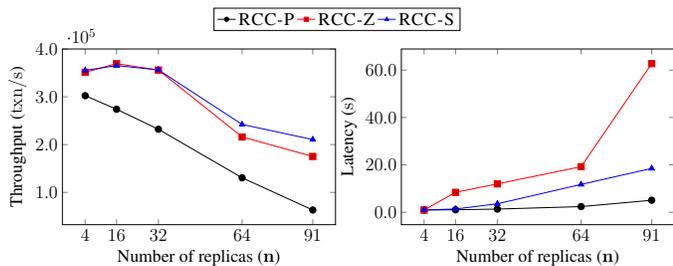


Fig. 9. Evaluating system throughput and latency attained by three RCC variants: RCC-P, RCC-Z and RCC-S when there are no failures.

### F. Analyzing RCC as a Paradigm

Finally, we experimentally illustrate the ability of RCC to act as a paradigm. To do so, we apply RCC to not only PBFT, but also to ZYZZYVA and SBFT. In Figure 9, we plot the performance of these three variants of RCC: RCC-P (RCC+PBFT), RCC-Z (RCC+ZYZZYVA), and RCC-S (RCC+SBFT). To evaluate the scalability of these protocols, we perform experiments in the optimistic setting with no failures and  $m = n$  concurrent instances.

It is evident from these plots that all RCC variants achieve extremely high throughput. As SBFT and ZYZZYVA only require linear communication in the optimistic case, RCC-S and RCC-Z are able to achieve up to  $3.33\times$  and  $2.78\times$  higher throughputs than RCC-P, respectively.

Notice that RCC-S consistently attains equal or higher throughput than RCC-Z, even though ZYZZYVA scales better than SBFT. This phenomena is caused by the way RCC-Z interacts with clients. In specific, like ZYZZYVA, RCC-Z requires its clients to wait for responses of all  $n$  replicas. Hence, clients have to wait longer to place new transactions, and consequently RCC-Z requires more clients than RCC-S to attain maximum performance. Even if we ran RCC-Z with 5 million clients, the largest amount at our disposal, we would not see maximum performance. Due to the low single-primary performance of ZYZZYVA, this phenomena does not prevent ZYZZYVA to already reach its maximum performance.

## VI. RELATED WORK

In Section I-A, we already discussed well-known primary-backup consensus protocols such as PBFT, ZYZZYVA, and HOTSTUFF and why these protocols are underutilizing resources. Furthermore, there is abundant literature on consensus and on primary-backup consensus in specific (e.g., [11], [33], [34], [35]). Next, we shall focus on the few works that deal with either *improving throughput and scalability* or with *improving resilience*, the two strengths of RCC

*Parallelizing consensus:* Several recent consensus designs propose to run several primaries concurrently, e.g., [20], [36], [37], [38]. None of these proposals satisfy all design goals of RCC, however. In specific, these proposals all fall short with respect to maximizing potential throughput in all cases, as none of these proposals satisfy the *wait-free* design goals D4 and D5 of RCC.

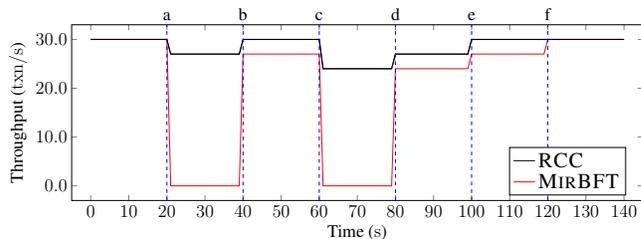


Fig. 10. Throughput of RCC versus MIRBFT during instance failures with  $m = 11$  instances. At (a), primary  $\mathcal{P}_1$  fails. In RCC, all other instances are unaffected, whereas in MIRBFT all replicas need to coordinate recovery. At (b), recovery is finished. In RCC, all instances can resume work, whereas MIRBFT halts an instance due to recovery. At (c) primaries  $\mathcal{P}_1$  and  $\mathcal{P}_2$  fail. In RCC,  $\mathcal{P}_2$  will be recovered at (d) and  $\mathcal{P}_1$  at (e) (as  $\mathcal{P}_1$  failed twice, its recovery in RCC takes twice as long). In MIRBFT, recovery is finished at (d), after which MIRBFT operates with only  $m = 9$  instances. At (e) and (f), MIRBFT decides that the system is sufficiently reliable, and MIRBFT enables the remaining instances one at a time.

**Example VI.1.** The MIRBFT protocol proposes to run concurrent instances of PBFT, this in a similar fashion as RCC. The key difference is how MIRBFT deals with failures: MIRBFT operates in global *epochs* in which a *super-primary* decides which instances are enabled. During any failure, MIRBFT will switch to a new epoch via a view-change protocol that temporarily shuts-down all instances and subsequently reduces throughput to zero. This is in sharp contrast to the wait-free design of RCC, in which failures are handled on a per-instance level. In Figure 10, we illustrated these differences in the failure recovery of RCC and MIRBFT.

As is clear from the figure, the fully-coordinated approach of MIRBFT results in substantial performance degradation during failure recovery. Hence, MIRBFT does not meet design goals D4 and D5, which is sharply limits the throughput of MIRBFT when compared to RCC.

*Reducing malicious behavior:* Several works have observed that traditional consensus protocols only address a narrow set of malicious behavior, namely behavior that prevents any progress [20], [22], [23], [39]. Hence, several designs have been proposed to also address behavior that impedes performance without completely preventing progress. One such design is RBFT, which uses concurrent primaries *not* to improve performance—as we propose—but only to mitigate throttling attacks in a way similar to what we described in Section IV. In practice, the design of RBFT results in poor performance at high costs.

HOTSTUFF [18], SPINNING [22], and PRIME [23] all propose to minimize the influence of malicious primaries by replacing the primary every round. This would not incur the costs of RBFT, while still reducing—but not eliminating—the impact of faulty replicas to severely reduce throughput. Unfortunately, these protocols follow the design of primary-backup consensus protocols and, as discussed in Section II, these designs are unable to achieve throughputs close to those reached by a concurrent consensus such as RCC.

*Concurrent consensus via sharding:* Several recent works have proposed to speed up consensus-based systems by in-

corporating sharding, this either at the data level (e.g., [5], [7], [40], [41], [42]) or at the consensus level (e.g., [43]). In these approaches only a small subset of all replicas, those in a single shard, participate in the consensus on any given transaction, thereby reducing the costs to replicate this transaction and enabling concurrent transaction processing in independent shards. As such, sharded designs can promise huge scalability benefits for easily-sharded workloads. To do so, sharded designs utilize a weaker failure model than the fully-replicated model RCC uses, however. Consider, e.g., a sharded system with  $z$  shards of  $n = 3f + 1$  replicas each. In this setting, the system can only tolerate failure of up to  $f$  replicas in a single shard, whereas a fully-replicated system using  $z$  replicas could tolerate the failure of any choice of  $\lfloor (zn - 1)/3 \rfloor$  replicas. Furthermore, sharded designs typically operate consensus protocols such as PBFT in each shard to order local transactions, which opens the opportunity of concurrent consensus and RCC to achieve even higher performance in these designs.

## VII. CONCLUSION

In this paper, we proposed *concurrent consensus* as a major step toward enabling high-throughput and more scalable consensus-based database systems. We have shown that concurrent consensus is in theory able to achieve throughputs that primary-backup consensus systems are unable to achieve. To put the idea of concurrent consensus in practice, we proposed the RCC paradigm that can be used to make normal primary-backup consensus protocols *concurrent*. Furthermore, we showed that RCC is capable of making consensus-based systems more resilient to failures by sharply reducing the impact of faulty replicas on the throughput and operations of the system. We have also put the design of the RCC paradigm to the test by implementing it in RESILIENTDB, our high-performance resilient blockchain fabric, and comparing it with state-of-the-art primary-backup consensus protocols. Our experiments show that RCC is able to fulfill the promises of concurrent consensus, as it significantly outperforms other consensus protocols and provides better scalability. As such, we believe that RCC opens the door to the development of new high-throughput resilient database and federated transaction processing systems.

*Acknowledgements:* We would like to acknowledge Sajjad Rahnama and Patrick J. Liao for their help during the initial stages of this work.

## REFERENCES

- [1] S. Gupta, J. Hellings, and M. Sadoghi, "Brief announcement: Revisiting consensus protocols through wait-free parallelization," in *33rd International Symposium on Distributed Computing (DISC 2019)*, vol. 146. Schloss Dagstuhl, 2019, pp. 44:1–44:3.
- [2] M. Herlihy, "Blockchains from a distributed computing perspective," *Commun. ACM*, vol. 62, no. 2, pp. 78–85, 2019.
- [3] A. Narayanan and J. Clark, "Bitcoin's academic pedigree," *Commun. ACM*, vol. 60, no. 12, pp. 36–45, 2017.
- [4] S. Gupta, J. Hellings, and M. Sadoghi, *Fault-Tolerant Distributed Transactions on Blockchains*, ser. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2020, (to appear).
- [5] M. J. Amiri, D. Agrawal, and A. E. Abbadi, "CAPER: A cross-application permissioned blockchain," *Proc. VLDB Endow.*, vol. 12, no. 11, pp. 1385–1398, 2019.
- [6] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, and J. Yellick, "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in *Proceedings of the Thirteenth EuroSys Conference*. ACM, 2018, pp. 30:1–30:15.
- [7] M. El-Hindi, C. Binnig, A. Arasu, D. Kossmann, and R. Ramamurthy, "BlockchainDB: A shared database on blockchains," *Proc. VLDB Endow.*, vol. 12, no. 11, pp. 1597–1609, 2019.
- [8] S. Nathan, C. Govindarajan, A. Saraf, M. Sethi, and P. Jayachandran, "Blockchain meets database: Design and implementation of a blockchain relational database," *Proc. VLDB Endow.*, vol. 12, no. 11, pp. 1539–1552, 2019.
- [9] F. Nawab and M. Sadoghi, "Blockplane: A global-scale byzantizing middleware," in *35th International Conference on Data Engineering (ICDE)*. IEEE, 2019, pp. 124–135.
- [10] L. Lao, Z. Li, S. Hou, B. Xiao, S. Guo, and Y. Yang, "A survey of IoT applications in blockchain systems: Architecture, consensus, and traffic modeling," *ACM Comput. Surv.*, vol. 53, no. 1, 2020.
- [11] C. Cachin and M. Vukolic, "Blockchain consensus protocols in the wild (keynote talk)," in *31st International Symposium on Distributed Computing*, vol. 91. Schloss Dagstuhl, 2017, pp. 1:1–1:16.
- [12] J. Gray, "Notes on data base operating systems," in *Operating Systems, An Advanced Course*. Springer-Verlag, 1978, pp. 393–481.
- [13] D. Skeen, "A quorum-based commit protocol," Cornell University, Tech. Rep., 1982.
- [14] S. Gupta and M. Sadoghi, "EasyCommit: A non-blocking two-phase commit protocol," in *Proceedings of the 21st International Conference on Extending Database Technology*. Open Proceedings, 2018, pp. 157–168.
- [15] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, 2002.
- [16] R. Kotla, L. Alvisi, M. Dahlin, A. Clement, and E. Wong, "Zyzyva: Speculative byzantine fault tolerance," *ACM Trans. Comput. Syst.*, vol. 27, no. 4, pp. 7:1–7:39, 2009.
- [17] G. Golan Gueta, I. Abraham, S. Grossman, D. Malkhi, B. Pinkas, M. Reiter, D.-A. Seredinschi, O. Tamir, and A. Tomescu, "SBFT: A scalable and decentralized trust infrastructure," in *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2019, pp. 568–580.
- [18] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hot-Stuff: BFT consensus with linearity and responsiveness," in *Proceedings of the ACM Symposium on Principles of Distributed Computing*. ACM, 2019, pp. 347–356.
- [19] S. Gupta, J. Hellings, S. Rahnama, and M. Sadoghi, "Proof-of-execution: Reaching consensus through fault-tolerant speculation," 2019. [Online]. Available: <http://arxiv.org/abs/1911.00838>
- [20] P.-L. Aublin, S. B. Mokhtar, and V. Quéma, "RBFT: Redundant byzantine fault tolerance," in *2013 IEEE 33rd International Conference on Distributed Computing Systems*. IEEE, 2013, pp. 297–306.
- [21] S. Gupta, S. Rahnama, and M. Sadoghi, "Permissioned blockchain through the looking glass: Architectural and implementation lessons learned," in *40th International Conference on Distributed Computing Systems*. IEEE, 2020.
- [22] G. S. Veronese, M. Correia, A. N. Bessani, and L. C. Lung, "Spin one's wheels? byzantine fault tolerance with a spinning primary," in *2009 28th IEEE International Symposium on Reliable Distributed Systems*. IEEE, 2009, pp. 135–144.
- [23] Y. Amir, B. Coan, J. Kirsch, and J. Lane, "Prime: Byzantine replication under attack," *IEEE Trans. Depend. Secure Comput.*, vol. 8, no. 4, pp. 564–577, 2011.
- [24] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *J. ACM*, vol. 32, no. 2, pp. 374–382, 1985.
- [25] S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," *SIGACT News*, vol. 33, no. 2, pp. 51–59, 2002.
- [26] J. Hellings and M. Sadoghi, "Coordination-free byzantine replication with minimal communication costs," in *23rd International Conference on Database Theory (ICDT 2020)*, vol. 155. Schloss Dagstuhl, 2020, pp. 17:1–17:20.

- [27] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, "Benchmarking cloud serving systems with YCSB," in *Proceedings of the 1st ACM Symposium on Cloud Computing*. ACM, 2010, pp. 143–154.
- [28] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "BLOCKBENCH: A framework for analyzing private blockchains," in *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM, 2017, pp. 1085–1100.
- [29] S. Rahnema, S. Gupta, T. Qadah, J. Hellings, and M. Sadoghi, "Scalable, resilient and configurable permissioned blockchain fabric," *Proc. VLDB Endow.*, vol. 13, no. 12, pp. 2893–2896, 2020.
- [30] S. Gupta, J. Hellings, S. Rahnema, and M. Sadoghi, "An in-depth look of BFT consensus in blockchain: Challenges and opportunities," in *Proceedings of the 20th International Middleware Conference Tutorials, Middleware*. ACM, 2019, pp. 6–10.
- [31] —, "Blockchain consensus unraveled: virtues and limitations," in *Proceedings of the 14th ACM International Conference on Distributed and Event-based Systems*. ACM, 2020, pp. 218–221.
- [32] —, "Building high throughput permissioned blockchain fabrics: Challenges and opportunities," *Proc. VLDB Endow.*, vol. 13, no. 12, pp. 3441–3444, 2020.
- [33] C. Berger and H. P. Reiser, "Scaling byzantine consensus: A broad analysis," in *Proceedings of the 2nd Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers*. ACM, 2018, pp. 13–18.
- [34] T. T. A. Dinh, R. Liu, M. Zhang, G. Chen, B. C. Ooi, and J. Wang, "Untangling blockchain: A data processing view of blockchain systems," *IEEE Trans. Knowl. Data Eng.*, vol. 30, no. 7, pp. 1366–1385, 2018.
- [35] S. Gupta and M. Sadoghi, *Blockchain Transaction Processing*. Springer International Publishing, 2018, pp. 1–11.
- [36] C. Stathakopoulou, T. David, and M. Vukolic, "Mir-BFT: High-throughput BFT for blockchains," 2019. [Online]. Available: <http://arxiv.org/abs/1906.05552>
- [37] M. Eischer and T. Distler, "Scalable byzantine fault-tolerant state-machine replication on heterogeneous servers," *Computing*, vol. 101, pp. 97–118, 2019.
- [38] B. Li, W. Xu, M. Z. Abid, T. Distler, and R. Kapitza, "SAREK: Optimistic parallel ordering in byzantine fault tolerance," in *2016 12th European Dependable Computing Conference (EDCC)*. IEEE, 2016, pp. 77–88.
- [39] A. Clement, E. Wong, L. Alvisi, M. Dahlin, and M. Marchetti, "Making byzantine fault tolerant systems tolerate byzantine faults," in *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*. USENIX Association, 2009, pp. 153–168.
- [40] J. Hellings, D. P. Hughes, J. Primero, and M. Sadoghi, "Cerberus: Minimalistic multi-shard byzantine-resilient transaction processing," 2020. [Online]. Available: <https://arxiv.org/abs/2008.04450>
- [41] M. J. Amiri, D. Agrawal, and A. El Abbadi, "SharPer: Sharding permissioned blockchains over network clusters," 2019. [Online]. Available: <https://arxiv.org/abs/1910.00765v1>
- [42] H. Dang, T. T. A. Dinh, D. Loghin, E.-C. Chang, Q. Lin, and B. C. Ooi, "Towards scaling blockchain systems via sharding," in *Proceedings of the 2019 International Conference on Management of Data*. ACM, 2019, pp. 123–140.
- [43] S. Gupta, S. Rahnema, J. Hellings, and M. Sadoghi, "ResilientDB: Global scale resilient blockchain fabric," *Proc. VLDB Endow.*, vol. 13, no. 6, pp. 868–883, 2020.