

# Canopus: A Scalable and Massively Parallel Consensus Protocol

Sajjad Rizvi  
University of Waterloo  
sm3rizvi@uwaterloo.ca

Bernard Wong  
University of Waterloo  
bernard@uwaterloo.ca

Srinivasan Keshav  
University of Waterloo  
keshav@uwaterloo.ca

## ABSTRACT

Achieving consensus among a set of distributed entities (or *participants*) is a fundamental problem at the heart of many distributed systems. A critical problem with most consensus protocols is that they do not scale well. As the number of participants trying to achieve consensus increases, increasing network traffic can quickly overwhelm the network from topology-oblivious broadcasts, or a central coordinator for centralized consensus protocols. Thus, either achieving strong consensus is restricted to a handful of participants, or developers must resort to weaker models of consensus.

We propose Canopus, a highly-parallel consensus protocol that is ‘plug-compatible’ with ZooKeeper, which exploits modern data center network topology, parallelism, and consensus semantics to achieve scalability with respect to the number of participants and throughput (i.e., the number of key-value reads/writes per second). In our prototype implementation, compared to EPaxos and ZooKeeper, Canopus increases throughput by more than 4x and 16x respectively for read-heavy workloads.

## CCS CONCEPTS

• **Computer systems organization** → **Dependable and fault-tolerant systems and networks**; *Distributed architectures*;

## KEYWORDS

Consensus, Scalable, Fault Tolerance, Geo-Distributed Replication, Network-Aware Protocol

### ACM Reference Format:

Sajjad Rizvi, Bernard Wong, and Srinivasan Keshav. 2017. Canopus: A Scalable and Massively Parallel Consensus Protocol. In *CoNEXT '17: CoNEXT '17: The 13th International Conference on emerging Networking EXperiments and Technologies*, December 12–15, 2017, Incheon, Republic of Korea. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3143361.3143394>

## 1 INTRODUCTION

In the past few years, an important class of distributed applications has emerged that requires agreement on the entries of a replicated transaction log or ledger. Examples include geo-replicated database systems that support high-volume, conflict-free transaction processing [31], and private blockchains [1–3] that continuously

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CoNEXT '17, December 12–15, 2017, Incheon, Republic of Korea

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5422-6/17/12...\$15.00  
<https://doi.org/10.1145/3143361.3143394>

add records to a distributed ledger. Each entry consists of a unit of computation that can read and modify the states of the application<sup>1</sup>. Entries are processed atomically in the order that they appear in the log or ledger.

Supporting these applications requires a consensus protocol that can scale to hundreds of nodes spanning multiple datacenters. The protocol must also be able to handle large volumes of requests, and be efficient for write-intensive workloads with nearly uniform popularity distributions.

Most commonly used consensus protocols, such as Paxos [17] and Raft [32], rely on a centralized coordinator to service client requests and replicate state changes. This introduces unavoidable latency and concentrates both processing load and network traffic at the coordinator, limiting their scalability.

Protocols such as EPaxos [27], Mencius [24], and S-Paxos [8] address scalability by moving from a single coordinator to a set of coordinators. However, as we will discuss in more detail in Section 2, message dissemination in these protocols is network topology-unaware. Their scalability is therefore still limited in wide-area deployments with restricted network link capacities. Paxos Quorum Leases [29] can improve the scalability of consensus protocols by assigning leases to non-coordinators, and serving read requests locally by lease holders. However, they are not well suited for the workloads required by our motivating applications.

In this paper, we introduce *Canopus*, a parallel, network-aware, and decentralized consensus protocol designed for large-scale deployments. It achieves scalability by (a) exploiting high performance broadcast support in modern switches and (b) parallelizing communication along an *overlay tree* [6]. Moreover, observing that modern datacenter networks use hardware redundancy, making them immune to single points of failure, and that they offer numerous disjoint paths between all pairs of servers, making network partitioning extremely rare [21, 22, 39], Canopus is optimized for the common case that an entire rack of servers never fails, and that the network is never partitioned. This greatly simplifies our design. We make three contributions:

- We present Canopus, a scalable and parallel consensus protocol.
- We provide a prototype that can be integrated with ZooKeeper to simplify deployment.
- We compare the performance of our prototype with both ZooKeeper and EPaxos. We find that our globally-distributed prototype implementation can handle up to 5 million transactions per second, and that, unlike EPaxos, the throughput increases with the number of servers participating in the consensus protocol.

<sup>1</sup>An equivalent implementation would be for each entry to store a reference and the parameters to a stored procedure.

## 2 RELATED WORK

Many systems have been designed to solve the consensus problem. We first discuss some well-known centralized and decentralized consensus protocols. We then discuss systems that optimize performance by exploiting specialized hardware and networking technologies. We also outline related work in the area of group communication protocols and discuss the role of consensus protocols in the context of private blockchains.

### 2.1 Centralized Coordination

Paxos [17] is the most widely used consensus protocol. Paxos participants elect a central coordinator to enforce consensus on a single value, re-electing the coordinator if it fails. Many variants of Paxos have been proposed over the years. In Fast Paxos [19], clients send consensus requests directly to the *acceptors* to reduce the number of rounds of message exchanges. In Generalized Paxos [18], nodes can execute commutative requests out-of-order. FGCC [40] builds upon Generalized Paxos by reducing the message complexity to resolve conflicts. Although these variants aim to improve different aspects of Paxos performance, none of them significantly improve the scalability of the protocol.

Ring-Paxos [26] and Multi-Ring Paxos [25] are atomic multicast protocols developed on top of Paxos. In these protocols, nodes organize into one or more rings, and the Paxos coordinator and acceptor roles are assigned to the nodes along the ring. The major drawback of these multicast protocols is that the latency of operations increases proportionally to the number of participants.

Several other coordinator-based consensus or atomic broadcast algorithms have been proposed that have the same scalability limitations as Paxos. These include Raft [32], a consensus algorithm with understandability as its main design goal, and Zab [16], an atomic broadcast protocol used in ZooKeeper [14].

### 2.2 Decentralized Coordination

EPaxos [27], Mencius [24], and S-Paxos [8] have been proposed to improve the scalability of consensus algorithms. These protocols follow a decentralized approach. For example, in EPaxos and Mencius, each node serves as the coordinator for the clients connected to that node. The nodes then forward requests to each other for conflict resolution and ordering. S-Paxos distributes the message dissemination load across nodes in a similar way. However, the nodes send the request IDs to a centralized coordinator that runs the Paxos algorithm to order the requests.

Two primary limitations restrict the scalability of these systems. First, they are not network topology-aware. For example, each node broadcasts requests to every other participating node, or at least to a quorum of the nodes. Thus, multiple instances of the same request and response may traverse one or more oversubscribed intra-data-center or high-latency wide-area links, which limits throughput. In contrast, Canopus organizes the nodes into a network-aware overlay tree, and remote results are retrieved *once* then shared among peers. Second, these protocols broadcast *both* read and write requests. Canopus does not broadcast read requests. Instead, Canopus introduces small delays to the read operations that ensure that read operations are correctly linearized with respect to any concurrent writes operations.

The Giraffe consensus protocol logically organizes nodes into an interior-node-disjoint forest [38]. It uses a variant of Paxos to order write requests within a tree and across trees in the forest. Unlike Canopus, Giraffe does not provide linearizable consistency, and each tree in Giraffe has its own coordinator, which limits its scalability.

Canopus is similar to the recently proposed AllConcur [33] consensus algorithm, which organizes nodes in the form of an overlay digraph. In each consensus round, each AllConcur node atomically broadcasts a message to its successors in the graph while keeping track of node failures. AllConcur nodes are expected to be part of the same InfiniBand network, typically within a single datacenter. In comparison, Canopus is a network-aware protocol that organizes the nodes in the form of an wide-area overlay tree. Nodes located in the same rack form a virtual group, which can use any reliable broadcast or stronger protocol to reach agreement. Canopus efficiently stitches together these virtual groups, as described later in this paper.

### 2.3 Consensus Exploiting Network Hardware

Several recent systems [9, 10, 15, 20, 34] exploit fast networking devices, e.g., low latency switches, or specialized hardware, such as FPGAs, to achieve consensus with high throughput and low latency. Since they require extensive in-network support, these systems can only be deployed within the same rack or the same datacenter. In contrast, Canopus is designed to provide consensus between nodes across globally-distributed datacenters. Network-hardware-aware consensus protocols are compatible with Canopus, in that these systems can be used to achieve consensus among nodes within a Canopus virtual group.

### 2.4 Group Communication Protocols

Many overlay and group communication protocols have been proposed in the research literature [23, 35]. These protocols efficiently and reliably disseminate messages across a large number of participants. Canopus uses the LOT group communication protocol [6], which is similar to Astrolabe [41], to disseminate protocol state. LOT embodies three insights to achieve high performance: (1) communication should be *topology aware*; most communication should occur between topologically-close nodes, (2) communication should be *hierarchical*, so that message dissemination to  $n$  nodes takes  $O(\log(n))$  steps, and (3) communication should be *parallel*, so that all nodes can make progress independent of each other.

### 2.5 Private Blockchains

Private blockchains [1–3] can be viewed as a distributed append-only ledger that is replicated across a large number of nodes. Blockchains can be used for many different applications such as cryptocurrencies, stocks settlements, fast currency conversions, and smart contracts. A common problem with blockchains is poor scalability, which is primarily due to the use of proof-of-work or topologically-unaware consensus protocols. Blockchains use consensus to resolve conflicts among concurrent append requests to the next ledger index. However, the throughputs of these protocols peak at a few tens of thousands of transactions/second with

tens of participants [42]. Canopus can be used as a scalable consensus protocol for private blockchains to significantly increase their throughput. Although the current design of Canopus does not provide Byzantine fault tolerance, we are integrating Byzantine fault-tolerance into Canopus in ongoing work.

## 2.6 Scalability of Read Requests

Several read optimizations have been proposed to reduce the latency of read requests for read-intensive workloads. Paxos Quorum Leases [29] (PQL) grants read leases to a set of objects to nodes that frequently access them. Lease holders can then serve read requests locally as long as the lease is valid. PQL performs well for workloads where data popularity is highly skewed. In Megastore [7], nodes can read data locally if the local-replica is the most up-to-date. Otherwise, nodes read from a majority of replicas. Canopus include an optional read optimization (§7.2) that allows the nodes to always read from the local replica.

## 3 SYSTEM ASSUMPTIONS

We now specify our six design assumptions:

**Network topology:** Canopus makes only two minor assumptions about the nature of its data center network. These are likely to hold for the vast majority of modern data center networks:

- 1. All machines running Canopus in the same data-center are hosted in the same rack, and are dual-connected to Top-of-Rack (ToR) switches** that provide low-latency connectivity to all the machines in its rack. If the number of Canopus nodes in a datacenter exceeds what can be hosted in a single rack, we assume that we have control over the physical location of the Canopus' nodes within a datacenter, including the location in the racks or whether Canopus runs in a VM or a physical machine.
- 2. ToR switches are connected to each other through a hierarchical network fabric.** In general, we expect latency to increase and bandwidth to decrease as hop count between the pair of communicating nodes increases.

Although Canopus will perform correctly (i.e., either achieve consensus or stall) without the next two assumptions, for acceptable performance, we also assume that:

- 3. Full rack failures are rare.** More precisely, we assume that although individual machines, ToR switches, and/or links may fail, it is rarely the case that all machines in a rack are simultaneously unreachable from off-rack machines. Single-failure fault tolerance is easily achieved by using dual-NIC machines connected to two ToR switches on each rack. In the rare cases of rack failures, Canopus halts until the racks have recovered. Moreover, rack failures do not cause the system to enter into an unrecoverable state. This is because the nodes persistently store their protocol state, and our system can resume when enough nodes in the failed racks recover.
- 4. Network partitions are rare.** A network partition within a rack requires both ToR switches to simultaneously fail, which we assume is rare. Similarly, we assume that the network provides multiple disjoint paths between ToR switches, so that the ToR switches are mutually reachable

despite a link or switch failure. Finally, we assume that there are multiple redundant links between data centers, so that a wide-area link failure does not cause the network to partition.

Many modern datacenters use network topologies that meet these two assumptions [11, 39], including Clos-like topologies such as fat-tree [5], VL2 [13], and leaf/spine networks. These failure assumptions simplify the design of Canopus and allow us to optimize Canopus for the common case.

In the unlikely event of a network partition, we believe it is reasonable for a consensus protocol to stall while waiting for the network to recover, resuming seamlessly once the network recovers. Indeed, this is the only reasonable response to a network partition where a majority of consensus participants do not reside in any single partition. We believe that, given both the rarity of network partitions and the severity in which they affect applications and services, it is not worthwhile for a consensus protocol to provide limited availability at the expense of a significant increase in protocol complexity. Planned shutdown of a rack is possible and does not result in any loss in availability.

**5. Client-coordinator communication:** In Canopus, as in many other consensus systems, clients send key-value read and write requests to a Canopus *node*. We assume that a client that has a pending request with a particular node does not send a request to any other node. This is to allow serialization of client requests at each node. If a client were to concurrently send a write and a read request to two different nodes, the read request could return before the write request even if the read was sent after the write, which would violate program order.

**6. Crash-stop failures:** We assume that nodes fail by crashing and require a failed node to rejoin the system using a join protocol. The Canopus join protocol closely resembles the join protocol for other consensus systems [32]. Canopus detects node failures by using a method similar to the heartbeat mechanism in Raft.

## 4 CANOPUS

Canopus is a scalable distributed consensus protocol that ensures live nodes in a Canopus group agree on the same ordered sequence of operations. Unlike most previous consensus protocols, Canopus does not have a single leader and uses a virtual tree overlay for message dissemination to limit network traffic across oversubscribed links. It leverages hardware redundancies, both within a rack and inside the network fabric, to reduce both protocol complexity and communication overhead. These design decisions enable Canopus to support large deployments without significant performance penalties.

The Canopus protocol divides execution into a sequence of *consensus cycles*. Each cycle is labelled with a monotonically increasing cycle ID. During a consensus cycle, the protocol determines the order of pending write requests received by nodes from clients before the start of the cycle and performs the write requests in the same order at every node in the group. Read requests are responded to by the node receiving it. Section 5 will describe how Canopus provides linearizable consistency while allowing any node to service read requests and without needing to disseminate read requests.

Canopus determines the ordering of the write requests by having each node, for each cycle, independently choose a large random number, then ordering write requests based on the random numbers. Ties are expected to be rare and are broken deterministically using the unique IDs of the nodes. Requests received by the same node are ordered by their order of arrival, which maintains request order for a client that sends multiple outstanding requests during the same consensus cycle.

## 4.1 Leaf-Only Trees

During each consensus cycle, each Canopus node disseminates the write requests it receives during the previous cycle to every other node in a series of *rounds*. For now, we will assume that all nodes start the same cycle in parallel at approximately the same time. We relax this assumption in Section 4.4, where we discuss how Canopus nodes self-synchronize.

Instead of directly broadcasting requests to every node in the group, which can create significant strain on oversubscribed links in a datacenter network or wide-area links in a multi-datacenter deployment, message dissemination follows paths on a topology-aware virtual tree overlay.

Specifically, Canopus uses a Leaf-Only Tree overlay [6], that allows nodes arranged in a logical tree to compute an arbitrary global aggregation function. We adapt LOT for use in a consensus protocol; from this point on, we will always be referring to our modified version of LOT.

A LOT has three distinguishing properties:

**i. Physical and virtual nodes:** In LOT, only the leaf-nodes exist physically (in the form of a dedicated process running in a physical machine). Internal nodes are virtual and do not exist as a dedicated process. Instead, each leaf node emulates *all of its ancestor nodes*. For clarity, we denote a leaf node as a *pnode* and an internal node as a *vnode*.

**ii. Node emulation:** Each pnode emulates all of its ancestor vnodes. i.e., each pnode is aware of and maintains the state corresponding to all of its ancestor vnodes. Thus, the current state of a vnode can be obtained by querying any one of its descendants, making vnodes inherently fault tolerant, and making vnode state access parallelizable.

**iii. Super-leaves:** All the pnodes located within the same rack are grouped into a single *super-leaf* for two reasons. First, this reduces the number of messages exchanged between any two super-leaves; instead of all-to-all communication between all the pnodes in the respective super-leaves, only a subset of the super-leaf nodes, called its *representatives*, communicate with another super-leaf on behalf of their peers. Second, because all the pnodes in a super-leaf replicate their common parents' state, a majority of the super-leaf members need to simultaneously fail to cause the super-leaf to fail.

Figure 1 shows an example of a leaf-only tree consisting of 27 pnodes. There are three pnodes per super-leaf. The pnode  $N$  emulates all of its ancestor vnodes 1.1.1, 1.1, and 1. The root node 1 is emulated by all of the pnodes in the tree.

We assume that during Canopus initialization, each node is provided with the identities and IP addresses of the peers in its own super-leaf.

## 4.2 Consensus Cycle

Each consensus cycle consists of  $h$  rounds where  $h$  is the height of the LOT. Each consensus cycle is labelled by a monotone non-decreasing *cycle ID*. Canopus ensures that nodes concurrently execute the same consensus cycle, which we further explain in Section 4.4. In the following sections, we describe the details of the messages exchanged in each round of a consensus cycle.


**First round in a consensus cycle:** In the first round of a consensus cycle, each pnode prepares a *proposal message* and broadcasts the message to the peers in its super-leaf using a reliable broadcast protocol, which we describe in detail in Section 4.3. A proposal message contains the set of pending client requests, group membership updates, and a large random number generated by each node at the start of the consensus cycle. This random number, called the *proposal number*, is used to order the proposals across the nodes in each round in the consensus cycle. A proposal message is also tagged with the cycle ID, the round number, and the vnode ID whose state is being represented by the proposal message. The group membership updates include the membership changes that happened before starting the current consensus cycle. Note that proposal messages sent in the first round are used to inform the receivers about the start of the next consensus cycle.

After receiving the round-1 proposal messages from all the peers in its super-leaf, each pnode independently orders the proposals according to their (random) proposal numbers. The sorted list of proposals represents the partial order of the requests in the current round and constitutes the state of the parent vnode of the super-leaf. A node finishes its first round after it has obtained proposal messages from all its super-leaf peers, and then immediately starts the next round in the consensus cycle.

**Round- $i$  in a consensus cycle:** In each subsequent round, each node independently prepares a new proposal message, which is used to share the computed state from the previous round with other super-leaves. The list of requests in the proposal message is simply the sorted list of requests obtained in the previous round. Similarly, the proposal number in the new message is the largest proposal number from the previous round. Thus, a proposal message for round  $i$  can be defined as  $M_i = \{R'_{i-1}, N'_{i-1}, F'_{i-1}, C, i, v\}$ , where  $R'$  is the sorted list of proposals from the previous round,  $N'_{i-1}$  is the largest proposal number in the previous round,  $F'_{i-1}$  is the set of membership updates received in the proposals from the previous round,  $C$  is the cycle ID,  $i$  is the round number, and  $v$  is the vnode ID, or the pnode ID in the case of the first round.

In round  $i$ , the pnodes compute, in parallel, the state of their height- $i$  ancestor vnode. For example, the root vnode in Figure 1 has a height of 3, and therefore, its state is computed in the third round. As the current state of a vnode is just the merged and sorted list of proposals belonging to its child pnodes or vnodes, each super-leaf independently gathers the proposals of that vnode's children from their emulators, where a vnode's emulator is a pnode that is known to have the vnode's state; see § 4.6 for details.

To obtain the proposal belonging to a vnode, representatives from each super-leaf select one of that vnode's emulators and sends a proposal-request message to the selected emulator. For instance, to compute the state of the node 1.1, representative node  $N$  in Figure 1 sends a proposal-request to emulators of the vnodes



**Figure 1: An example of a leaf-only tree (LOT). Only the leaf nodes exist physically and the internal nodes are virtual. A leaf node emulates all of its ancestor virtual nodes.**

1.1.2 and 1.1.3 – it already knows the state of the vnode 1.1.1, which was computed in the previous round.

The proposal-request message contains the cycle ID, round number, and the vnode ID for which the state is required. In response to receiving a proposal-request message, the receiver pnode replies with the proposal message belonging to the required vnode. If the receiver pnode has not yet computed the state of the vnode, it buffers the request message and replies with the proposal message only after computing the state of the vnode. When the super-leaf representative receives the proposal message, it broadcasts the proposal to its super-leaf peers using reliable broadcast. These peers then independently sort the proposals, resulting in each node computing the state of their height- $i$  ancestor. Note that representatives from all super-leaves issue proposal-requests in parallel, and all nodes in all super-leaves independently and in parallel compute the state of their height- $i$  ancestor.

**Completing a consensus cycle:** The consensus protocol finishes after all nodes have computed the state of the root node, which takes time proportional to the logarithm of the number of Canopus pnodes. At this point, every node has a total order of requests received by all the nodes. The node logs the requests in its request log for execution. It then initiates the next consensus cycle with its super-leaf peers if it received one or more client requests during the prior consensus cycle.

### 4.3 Reliable-Broadcast in a Super-Leaf

The nodes in a super-leaf use reliable broadcast to exchange proposal messages with each other. For ToR switches that support hardware-assisted atomic broadcast, nodes in a super-leaf can use this functionality to efficiently and safely distribute proposal messages within the super-leaf.

If hardware support is not available, we use a variant of Raft [32] to perform reliable broadcast. Each node in a super-leaf creates its own dedicated Raft group and becomes the initial *leader* of the group. All other nodes in the super-leaf participate as *followers* in the group. Each node broadcasts its proposal messages to the other nodes in its super-leaf using the Raft log replication protocol in its own Raft group. If a node fails, the other nodes detect that the leader of the group has failed, and elect a new leader for the group

using the Raft election protocol. The new leader completes any incomplete log replication, after which all the nodes leave that group to eliminate the group from the super-leaf.

Reliable broadcast requires  $2F + 1$  replicas to support  $F$  failures. If more than  $F$  nodes fail in the same super-leaf, the entire super-leaf fails and the consensus process halts.

### 4.4 Self-Synchronization

So far we have assumed that all of the nodes start at the same consensus cycle in parallel. Here, we describe how the nodes are self-synchronized to concurrently execute the same consensus cycle.

Before starting a consensus cycle, a node remains in an idle state. In this state, its network traffic is limited to periodic heartbeat messages to maintain its liveness. A new consensus cycle only starts when the node receives outside prompting. The simplest version of this prompting is when the node receives a client request for consensus. A client request triggers the node to start a new consensus cycle and broadcast a proposal message to other peers in its super-leaf. The proposal message informs the peers that it is time to start the next consensus cycle, if it is not already under way. The representatives, in effect, send proposal-request messages to other super-leaves that trigger them to start the next consensus cycle, if they have not started yet.

Alternatively, a node, while in the idle state, may receive either a proposal message from its peers in its super-leaf, or a proposal-request message from another super-leaf. This indicates that a new cycle has started, and therefore, this node begins the next consensus cycle. In this scenario, the node's proposal message contains an empty list of client requests. As a result, the nodes are self-locked to execute the same consensus cycle.

### 4.5 Super-Leaf Representatives

A super-leaf representative fetches the state of a required vnode on behalf of its super-leaf peers and broadcasts the fetched state to the peers using reliable broadcast (§ 4.3). We use a variant of Raft's leader election and heartbeat protocols [32] to allow nodes in a super-leaf to select one or more super-leaf representatives and detect representative failures. A representative failure triggers a new election to select a replacement representative.

The representatives of a super-leaf are each individually responsible for fetching the states of the tree's vnodes. A representative can fetch the states of more than one vnode. However, to improve load balancing, we recommend that different representatives fetch the states of different vnodes. The vnode to representative assignment can be performed deterministically by taking the modulo of each vnode ID by the number of representatives in a super-leaf. Because representatives of a super-leaf are numbered and ordered (e.g., representative-0, representative-1, etc.), the result of the modulo operation can be used to determine which representative is assigned to a vnode without the need for any communication or consensus.

#### 4.6 Emulation Table


Thus far, we have not discussed how Canopus nodes actually communicate with other nodes. To initiate communication with an emulator of a vnode (a pnode in the sub-tree of the vnode) with a given ID, LOT maintains an *emulation table* that maps each vnode to a set of IP addresses of the pnodes that emulate that vnode. For example, for the LOT given in Figure 1, the table would map vnode 1.1 to the set of IP addresses of the nine pnodes that emulate this vnode.

We assume that at system initialization time, the emulation table is complete, correct, and made available to all pnodes. Subsequently, failed pnodes must be removed from the emulation tables maintained by all non-failed pnodes. Note that this itself requires consensus. Hence we maintain emulation tables *in parallel* with the consensus process, by piggybacking changes in the emulation table on proposal messages.


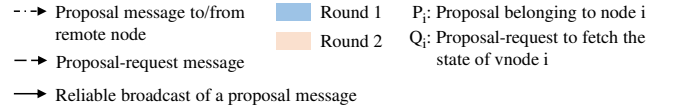
Specifically, in the first round of a consensus cycle, pnodes list membership changes (determined based on an within-super-leaf failure-detector protocol, such as through Raft heartbeats) in their proposal messages. Membership changes for a pnode include the list of pnodes that joined or left its super-leaf before the pnode started its current consensus cycle. At the end of the consensus cycle, all the pnodes have the same set of proposal messages, and thus the same set of membership changes are delivered to each pnode. Therefore, each pnode applies the same membership updates to the emulation table at the end of each consensus cycle. As a result, each pnode has the same emulation table and same membership view of LOT in each consensus cycle.

#### 4.7 An Illustrative Example

We illustrate Canopus using an example in which six nodes are arranged in two super-leaves in a LOT of height 2 as shown in Figure 2(a). Nodes A, B, and C comprise super-leaf  $S_x$ , and the nodes D, E, and F comprise the super-leaf  $S_y$ . The nodes in  $S_x$  and  $S_y$  are connected with the common vnodes  $x$  and  $y$  respectively. The vnodes  $x$  and  $y$  are connected with the root vnode  $z$ . In the first round, the states of vnodes  $x$  and  $y$  are computed, which establishes consensus within the super-leaves. In the second round, the state of the root vnode  $z$  is computed, which establishes consensus between super-leaves. Figure 2(b) shows the events occurring in the super-leaf  $S_x$  during the consensus cycle. The events are labelled in circles and explained as follows:



(a) Six nodes are arranged in two super-leaves in a LOT of height 2.



(b) Timing diagram of the events occurring in the super-leaf  $S_x$ , which consists of the nodes A, B, and C.

**Figure 2: Illustration of a consensus cycle in Canopus.**

**1. The consensus cycle starts:** Nodes A and B start the first round of the consensus cycle. At the start of the cycle, nodes A and B have the pending requests  $R_A$  and  $R_B$ . The nodes prepare proposal-messages and broadcast them to their peers in the super-leaf.

Assume that the proposals of nodes A, B, and C are  $P_A = \{R_A \mid N_A \mid 1\}$ ,  $P_B = \{R_B \mid N_B \mid 1\}$ , and  $P_C = \{\phi \mid N_C \mid 1\}$ , where  $R_i$  is the ordered set of pending requests on node  $i$ ,  $N_i$  is the proposal number, and 1 is the round number. We omit the other information from the proposals for simplicity.

**2. A proposal-request message is sent:** Node C receives the proposal message from node A and starts its consensus cycle. Assuming that nodes A and C are the representatives for super-leaf  $S_x$ , these nodes send proposal-request messages  $Q_y$  to D and E respectively to redundantly fetch the state of vnode  $y$ . The nodes in  $S_x$  require the state of  $y$  to compute the state of the ancestor vnode  $z$  in the next round.

**3. A proposal-request is received for an unfinished round:** Node A receives a proposal-request message from F, which belongs to the next round of the current consensus cycle, asking for the proposal  $P_x$ . As node A has not yet finished its first round that results in  $P_x$ , it buffers the request and does not send any reply until it finishes its first round.

**4. The first round is completed:** Node C receives the proposal message  $P_B$  from B. As node C has received the proposals from

all the peers in its super-leaf, it finishes its first round. Node  $C$  sorts the requests in the received proposals and its own proposal according to the proposal numbers. The sorted list of requests is the consensus result from the first round, and comprises the current state of the vnode  $x$ . When nodes  $A$  and  $B$  finish their first round, they also have the same sorted list of requests.

To continue with the next round, the nodes prepare the proposal-messages that contain the sorted list of requests and the largest proposal number from the last round. In this example, the three nodes prepare the round-2 proposal  $P_x = \{P_A, P_C, P_B \mid N_B \mid 2\}$ , assuming  $N_A < N_C < N_B$ .

**5. A proposal-request is served:** Node  $A$  continues in the next round after finishing the first round and prepares the proposal  $P_x$ . It replies to the pending proposal-request  $Q_x$ , which it received from node  $F$ .

**6. Coordinating the proposals:** Node  $C$  receives the reply of its request  $Q_y$  for the round-2 proposal  $P_y$ , and reliably broadcasts  $P_y$  to other peers in its super-leaf. Assume that  $P_y = \{P_D, P_E, P_F \mid N_F \mid 2\}$ .

**7. The consensus cycle is completed:** After receiving all the required round-2 proposals, node  $C$  finishes the second round of the consensus cycle. At the end of the round, node  $C$  has the proposal  $P_z = \{P_y, P_x \mid N_B \mid 3\}$ , assuming  $N_F < N_B$ . As node  $C$  has now finished calculating the state of the root node, it finishes the consensus cycle. At the end of the cycle, node  $C$  has the consensus on the requests, which is  $\{R_D, R_E, R_F, R_A, R_C, R_B\}$ . Other nodes will have the same ordering of the requests when their second round is finished. Node  $C$  applies the requests in its copy of the commit log. Note that all node eventually agree on identical copies of this log.

## 5 LINEARIZABILITY

Canopus provides linearizability by totally ordering both read and update requests. Interestingly, Canopus enforces global total order *without* disseminating read requests to participating nodes, significantly reducing network utilization for read-heavy workloads. Instead, Canopus nodes delay read requests until all concurrently-received update requests are ordered through the consensus process. Once the update request ordering is known, each Canopus node locally orders its pending read requests with respect to its own update requests such that the request orders of its clients are not violated. This trades off read latency for network bandwidth.

During consensus cycle  $C_i$ , the nodes buffer the requests received from the clients. The requests accumulated on a node  $N_i$  during cycle  $C_j$  form a request set  $S_i^j$ . In Canopus, the requests within a request set are sequential and follow the receive order. In contrast, the request sets across the nodes in cycle  $C_j$ , which we denote as  $S_{*}^j$ , are concurrent. To order the concurrent sets in  $S_{*}^j$ , Canopus executes the next instance  $C_{j+1}$  of the consensus protocol. The consensus process gives a total order of the concurrent sets in  $S_{*}^j$ . The total order of the concurrent sets translates to the linearized order of the requests received during the cycle  $C_j$ . Note that the requests in a request set are never separated. We are instead ordering the request sets to determine a total order of requests. At the end of the consensus process, each node has the same total order

of write requests, and inserts its read requests in the appropriate positions to preserve its own request set order.

In this approach, each node delays its read requests for either one or two consensus cycles to provide linearizability. If a request is received immediately after starting a consensus cycle  $C_j$ , the read request is delayed for two consensus cycles: The currently executing cycle  $C_j$ , and the next cycle  $C_{j+1}$  in which the request sets  $S_i^j$  are ordered and executed. If the request is received just before completing the currently executing cycle  $C_j$ , the read is delayed for only one consensus cycle  $C_{j+1}$ .

## 6 CORRECTNESS PROPERTIES

Similar to Paxos and EPaxos, Canopus provides the properties of nontriviality, agreement, FIFO order of client requests, linearizability, and liveness<sup>2</sup>. Nontriviality and FIFO ordering are met by design. The following theorem ensures agreement, linearizability, and liveness:

**THEOREM 1:** For a height- $h$  LOT consisting of  $n$  super-leaves, all the live descendant-nodes  $\hat{E}_w^r$  of the root-vnode  $w$  that eventually complete the last round  $r = h$  have the same ordered-set of messages.

$$\forall_{i, j \in \hat{E}_w^r} M_i^r = M_j^r \quad (1)$$

where  $M_i^r$  is the set of messages that a live node  $i$  has after completing the round  $r$  of the current consensus cycle.

We provide a proof sketch for the special case of the first consensus cycle for height-2 tree with  $n$  super-leaves. The complete proof involves induction first on the number of cycles and then on the tree height. The complete proof is available in our extended technical report [36].

The proof assumes that:

- A1 All nodes are initialized with the same emulation table and membership view.
- A2 The network cannot be partitioned, messages are not corrupted, messages are eventually delivered to a live receiver, and that nodes fail by crashing. Node failures within a super-leaf are detected by using heartbeats and timeouts similar to Raft.
- A3 The super-leaf-level structure of LOT does not change: Nodes may leave or join super-leaves, however, super-leaves are not added or removed.
- A4 Reliable broadcast functionality is available within a super-leaf, which ensures that all the live nodes in a super-leaf receive the same set of messages.

**Proof Sketch:** Consider a height-2 LOT consisting of  $n$  super-leaves, as shown in Figure 3. Each super-leaf  $S_i$  has a height-1 parent vnode  $u_i$ . The height-1 vnodes are connected with the common root vnode  $w$ . A consensus cycle consists of two rounds for a LOT of height 2. In the first round, the nodes share the pending batch of update-requests within their super-leaf using the reliable broadcast functionality (assumption A4). The nodes sort the messages according to the proposal numbers received in the messages. At the end of the first round, each node  $f_{S_j}^i$  in a super-leaf  $S_j$  has the

<sup>2</sup>Please refer to Reference [27] for definitions.




Figure 3: A LOT of height 2 with  $n$  super-leaves.

same ordered-set of messages  $M_j^1$ , which represents the state of the height-1 ancestor  $u_j$  of the nodes in  $S_j$ . To order the messages in the next round, the largest proposal number received in the current round is selected.

In the second round each node  $f_{S_j}^i$  calculates the state of its height-2 ancestor  $w$ . Therefore, the set of messages that  $f_{S_j}^i$  requires to complete the second round is  $M_j^2 = \bigvee_{u \in C_j^2} \cup M_u^1$ , where  $C_j^2$  is the set of children of the height-2 ancestor of any node in the super-leaf  $S_j$ .

By Canopus' design, a set of representatives of  $S_j$  fetch the missing state  $M_u^1$  of each vnode  $u \in C_j^2$  from the live descendants of  $u$ . By assumption A1 and A2, the representatives will eventually succeed in fetching the state of  $u$  if there is at least one live super-leaf descending from  $u$ . If the representatives of  $S_j$  eventually succeed in fetching the states of all the vnodes  $u \in C_j^2$ , then by reliable broadcast within the super-leaf (assumption A4), all the nodes in the  $S_j$  have the same set of messages  $M_j^2$  that are required to complete the second round.

By assumptions A1 and A3, all the nodes have the same view of LOT in the first consensus cycle. Therefore, the representatives of each super-leaf fetch the same state of a common vnode  $u$  from any of the descendants of  $u$ . This implies that, for all the super-leaves for which the representatives have succeeded in fetching all the required states of height-1 vnodes, the nodes in the super-leaves have same set of messages, which allows the nodes to complete the second round. Thus,

$$\bigvee_{i,j \in \mathcal{E}_w^2} M_i^2 = M_j^2 \quad (2)$$

the nodes sort the messages according the proposal numbers received with the vnode-states. As the consensus cycle consists of two rounds for a LOT of height 2, equation 2 implies that all of the descendants of the height-2 root vnode that complete the first consensus cycle have the same ordered-set of messages. This shows that the live nodes reach agreement and eventually complete the consensus cycle, if a super-leaf does not fail.

If the representatives of a super-leaf  $S_j$  fail in fetching the state of at least one vnode  $u \in C_j^2$  in the second round, then either (a) all the descendants of  $u$  have crashed or they do not have the state of  $u$  due to not completing the previous round, or (b) all the representatives of  $S_j$  have crashed, which implies that the super-leaf  $S_j$  has failed due to insufficient number of live nodes to elect a new

representative. In either case, the nodes in  $S_j$  cannot complete the second round due to missing a required state, and the consensus process stalls for the nodes in  $S_j$  in the second round of the current consensus cycle.

Alternatively, if the nodes in another super-leaf  $S_k$  succeed in fetching all the required states, then the consensus process stalls for the nodes in  $S_k$  in the next cycle  $c + 1$  because the live nodes in  $S_j$  are stalled in the last cycle  $c$ . This shows that the consensus process stalls for the live nodes due to a super-leaf failure and the nodes do not return a wrong result.

In the subsequent consensus cycles, the assumption A1 does not hold because the membership of LOT changes if a node fails or a new node joins a super-leaf. However, we show in the detailed proof, which is available on request, using induction that the emulation tables are maintained and all the live nodes that eventually complete a cycle  $c$  have the same emulation table in cycle  $c + 1$ . Therefore, Canopus satisfies agreement and liveness properties in any consensus cycle for height-2 LOT. We prove these properties for a LOT of any height- $h$  using induction, for which the height-2 LOT serves as the base case.

## 7 OPTIMIZATIONS

In this section, we discuss two optimizations that enable Canopus to achieve high throughput and low read latency in wide-area deployments.

### 7.1 Pipelining

The completion time of a consensus cycle depends on the latency between the most widely-separated super-leaves. If only one cycle executes at a time, then high latency would reduce throughput as the nodes are mostly idle and are waiting for distant messages to arrive. Therefore, we use pipelining in Canopus to increase the throughput of the system. Specifically, instead of executing only one consensus cycle at a time, a node is permitted to participate in the next consensus cycle if any of the following events occur:

- The node receives a message belonging to the consensus cycle that is higher than its most recently started consensus cycle.
- A periodical timer expires, which serves as an upper bound for the offset between the start of two consensus cycles.
- The number of outstanding client requests exceeds the maximum batch size of requests.

With this change, a node can start exchanging messages with its peers that belong to the next consensus cycle even before the end of the prior cycle. In effect, this allows Canopus to maintain multiple in-progress consensus cycles. However, log commits only happen when a particular consensus cycle is complete, and always in strict order of consensus cycles. This is no different than a transport layer maintaining a window of transmitted-but-unacked packets.

Two cases need to be addressed to maintain the total order of the requests. First, nodes must always start consensus cycles in sequence (i.e., they cannot skip a cycle). It may happen that a node receives a message belonging to cycle  $C_{j \geq i+2}$  where  $C_i$  is the most recently started consensus cycle at the node. In that case, the node



still starts the next cycle  $C_{i+1}$  instead of starting the cycle  $C_j$ . Second, nodes always commit the requests from consensus cycles in sequence. Due to parallelism and fetching proposals from different nodes in each cycle, it may happen that a node receives all the required messages to complete a cycle  $C_{j \geq i+1}$  before it has completed the cycle  $C_i$ . However, to maintain the total order, nodes do not commit the requests from  $C_j$  until the requests from all the cycles before  $C_j$  have been committed.

## 7.2 Optimizing Read Operations

Canopus delays read operations to linearize them relative to the write requests across all the nodes in the system. Therefore, the read operations must be deferred to the end of the next consensus cycle, which is bounded by the round-trip latency between the farthest super-leaves. However, for read-heavy workload, it is desirable that the read operations are performed with minimal latency. Canopus can optionally support write-leases for a particular key during a consensus cycle to reduce the latency of read operations while preserving linearizability.

For any key, during any consensus cycle either a write lease is inactive, so that *no* writes are permitted to that key, and all nodes can read this key with no loss of consistency or *all* nodes have permission to write to this key (with a write order that is decided at the end of the consensus cycle) and no reads are permitted to this key. Any read requests made to a key during a consensus cycle  $i$  such that there is a write lease for the key is deferred to the end of the  $(i + 1)^{th}$  consensus cycle, to ensure linearizability.

Write leases require the following three modifications to Canopus:

- 1. Blocking client requests:** Following the model in Paxos Quorum Leases [30], our read optimization restricts the clients to sending only *one* request at any one time. This is different than our earlier model, where the clients are permitted many outstanding requests at a time.

- 2. Write leases:** Write requests are committed after acquiring the write-lease for the keys in the requests. Nodes piggyback lease requests with the proposal messages in the next consensus cycle  $C_{i+1}$ . At the end of the consensus cycle  $C_{i+1}$ , all the correct nodes that complete the cycle have the same set of lease requests. Therefore, all the correct nodes in the system can apply the lease in the same consensus cycle  $C_{i+p+1}$ .

- 3. Reads without delay:** A node  $N$  performs a read operation for a key  $y$  immediately (reading results from committed consensus cycles) if the write-lease for  $y$  is not active in any of the currently ongoing consensus cycles. If a node receives a read request for  $y$  while the write-lease for  $y$  is currently active, the node delays the read request until the end of the next consensus cycle.

## 8 EVALUATION

The focus of our evaluation is to compare the throughput and latency of Canopus with other competing systems at different deployment sizes. Our experiments are based on our prototype implementation of Canopus, which consists of approximately 2000 lines of Java code. We select EPaxos as a representative for state-of-the-art decentralized consensus approaches as EPaxos has been reported [27] to perform equal or better than other decentralized

consensus protocols. We use the publicly available implementation [28] of EPaxos from the paper’s authors for the evaluation. We also compare against ZooKeeper, a widely used coordination service that uses Zab, a centralized atomic broadcast protocol, to provide agreement. In order to provide a fair comparison with ZooKeeper, which includes additional abstractions that may introduce extra overhead, we created ZKCanopus, a modified version of ZooKeeper that replaces Zab with Canopus and performed all of our comparisons to ZooKeeper using ZKCanopus.

## 8.1 Single Datacenter Deployment

In this section, we evaluate the performance of Canopus and EPaxos when all the nodes are located in the same datacenter.

**Experimental setup:** We conducted our experiments on a three rack cluster where each rack consists of 13 machines. Each machine contains 32GB of memory, a 200GB Intel S3700 SSD, and two Intel Xeon E5-2620 processors. The machines in each rack are connected through a Mellanox SX1012 top-of-rack switch via 10Gbps links. Rack switches are connected through 2x10Gbps links to a common Mellanox SX1012 aggregation switch. Thus, the network oversubscription is 1.5, 2.5, 3.5, and 4.5 for experiments with 9, 15, 21, and 27 nodes respectively.


For our ZooKeeper and ZKCanopus experiments, both systems were configured to write logs and snapshots of the current system state asynchronously to the filesystem. We use an in-memory filesystem to simplify our experimental setup. To ensure that using an in-memory filesystem does not appreciably affect our results, we perform additional experiments in which logs are stored to an SSD. The results show that, for both ZooKeeper and ZKCanopus, throughput is not affected and median completion time increases by less than 0.5 ms.

Each experiment runs for 50 seconds and is repeated five times. We discard the first and last 5 seconds to capture the steady state performance. The error bars in the graphs show 95% confidence intervals.

**Workload:** Our experimental workloads are driven by 180 clients, which are uniformly distributed across 15 dedicated machines with 5 machines per rack. Each client connects to a uniformly-selected node in the same rack. The clients send requests to nodes according to a Poisson process at a given inter-arrival rate. Each request consists of a 16-byte key-value pair where the key is randomly selected from 1 million keys.

**LOT configuration in Canopus:** Canopus nodes are organized into three super-leaves in a LOT of height 2. To vary the group size, we change the super-leaf size to 3, 5, 7, and 9 nodes. Due to the size of our cluster, we did not run experiments with taller trees. In our current implementation, nodes reliably broadcast messages within a super-leaf without hardware assistance. Instead, they use the Raft-based reliable broadcast protocol described in Section 4.3 in which nodes communicate using unicast messages.

**Performance metrics:** The two key performance metrics we measure in our experiments are throughput and request completion time. We determine the throughput of a system by increasing the request inter-arrival rate until the throughput reaches a plateau. Our empirical results show that, for all of the tested systems and



**Figure 4: Throughput and median request completion times of Canopus and EPaxos while scaling the number of nodes in the system.**

configurations, the plateau is reached before the median request completion time reaches 10 ms. To simplify our testing procedure, our experiments run until the request completion time is above 10 ms and we use the last data point as the throughput result for that run. For each run, we manually verify that we have reached the system’s maximum throughput.

As it is difficult to determine the exact inter-arrival rate at which maximum throughput has been reached, together with the common operating practice of running critical systems below their maximum load, we report the median request completion time of the tested systems when they are operating at 70% of their maximum throughput. We believe this is more representative of the completion times that applications will experience when interacting with these systems compared to the median completion time at 100% load.

**8.1.1 Comparison with EPaxos.** In this experiment, we compare the performance of Canopus with EPaxos at different system sizes. Figure 4(a) shows the maximum throughput of the two systems. The first three bars in each group show the Canopus’ throughput at 20%, 50%, and 100% write requests, and the last two bars show the throughput of EPaxos at its default batch duration of 5 ms, in which requests are delayed for up to 5 ms in order to collect a larger batch of requests, and a reduced batch duration of 2 ms. For EPaxos, we ensure there is 0% command interference to evaluate its performance under the best possible condition. As EPaxos sends reads over the network to other nodes, its performance is largely independent of the write percentage; we show its results using the 20% write request workload.

The performance results show that with 20% and 50% write requests, Canopus provides significantly higher throughput than EPaxos with the performance gap increasing with larger system sizes. With 27 nodes and 20% write requests, Canopus provides more than 3 times the throughput of EPaxos with 5 ms batching. This is because Canopus can serve read requests locally without disseminating the request to the rest of the nodes in the group while still providing linearizable consistency. For the 100% write request workload, Canopus provides similar throughput to EPaxos

with a 5 ms batching duration. This is in part due to network bandwidth being relatively abundant in a single datacenter deployment. EPaxos with a 2 ms batching duration shows a significant drop in throughput as we increase the number of nodes in the system. This illustrates that EPaxos has scalability limitations when configured with smaller batch sizes.

Figure 4(b) shows the median request completion times of the two systems. From 9 to 27 nodes, Canopus’ request completion time is mostly independent of write request percentage, and is significantly shorter than EPaxos.

EPaxos’ high request completion time is primarily due to its batching duration. By reducing its batching duration to 2 ms, EPaxos’ median request completion time reduces by more than half and is within 1 ms of Canopus’ request completion time. However, as we saw previously, EPaxos relies on large batch sizes to scale. Therefore, our results show that EPaxos must tradeoff request completion time for scalability, whereas Canopus’ read throughput increases with more nodes, its write throughput remains largely constant up to at least 27 nodes, and its median request completion time only marginally increases going from 9 to 27 nodes.

**8.1.2 Comparison with ZooKeeper.** In this experiment, we compare the performance of ZooKeeper with ZKCanopus, our Canopus-integrated version of ZooKeeper. This comparison serves three purposes. First, it compares the performance of Canopus with that of a centralized coordinator-based consensus protocol. Second, it shows the effectiveness of Canopus in scaling an existing system by eliminating the centralized coordinator bottleneck. Finally, these experiments evaluate the end-to-end performance of complete coordination services, instead of evaluating just the performance of the consensus protocols.

We configure ZooKeeper to have only five *followers* with the rest of the nodes set as *observers* that do not participate in the Zab atomic broadcast protocol but asynchronously receive update requests. This choice is mainly to reduce the load on the centralized *leader* node. Although observers do not provide additional fault

**Table 1: Latencies (ms) between the datacenters used in the experiments.**

	IR	CA	VA	TK	OR	SY	FF
IR	0.2						
CA	133	0.2					
VA	66	60	0.25				
TK	243	113	145	0.13			
OR	154	20	80	100	0.26		
SY	295	168	226	103	161	0.2	
FF	22	145	89	226	156	322	0.23

tolerance, they can improve read request performance by servicing read requests. In the case of ZKCanopus, every node participates fully in the Canopus consensus protocol. We use one *znode* in these experiments and the clients read from and write to the same *znode*.

Figure 5 shows the throughput to median request completion time results for ZKCanopus and ZooKeeper with 9 and 27 nodes. For ZKCanopus, we do not show the request completion time at low throughputs in order to improve the readability of the graphs. The results show that, when the system is not loaded, ZKCanopus has a median request completion time that is between 0.2 ms to 0.5 ms higher than ZooKeeper. This is in part due to the additional network round trips required to disseminate write requests using a tree overlay instead of through direct broadcast. However, we believe the small increase in request completion time is acceptable given the significant improvement in scalability and throughput that ZKCanopus provides over ZooKeeper.

## 8.2 Multi-Datacenter Deployment

In this section, we evaluate the performance of Canopus in a wide area deployment. We run the experiments on Amazon EC2 cloud using three, five, and seven datacenters. Each datacenter consists of three nodes located in the same site. Each node runs on an EC2 c3.4xlarge instance, which consists of 30GB memory and Intel Xeon E5-2680 processor. Each datacenter has 100 clients that connect to a uniform-randomly selected node in the same datacenter and concurrently perform read and write operations at a given rate. The workload consists of 20% write requests, unless otherwise mentioned. The messages consists of 16-byte key-value pairs. The inter-datacenter latencies are given in the Table 1.

For Canopus, each datacenter contains a super-leaf, although the nodes might not be located in the same rack. We enable pipelining to overcome the long delays across datacenters. Each node starts a new consensus cycle every 5ms or after 1000 requests have accumulated, whichever happens earlier. For EPaxos, we used the same batch size as Canopus, and the workload consists of zero command interference. We enable latency-probing in EPaxos to dynamically select the nodes in the quorum. We disable thrifty optimization in EPaxos because our experiments show lower throughput with the thrifty optimization turned on.

Figure 6 shows the throughput and median request completion times of Canopus and EPaxos when deployed in three, five, and seven datacenters. Each datacenter consists of three nodes. The

vertical lines in the graph show the throughput when the latency touches 1.5 times the base latency.

Scaling across the datacenters, Canopus is able to achieve about 2.6, 3.8, and 4.7 millions requests per second, which is nearly 4x to 13.6x higher throughput than EPaxos. Canopus is able to achieve high throughput because of its more efficient utilization of CPU and network resources. Unlike EPaxos, Canopus does not disseminate read requests across the nodes, not even the read keys, which reduces the number and size of messages exchanged between the nodes. For ready-heavy workloads, this significantly impacts the utilization of CPU and network resources. Furthermore, due to the hierarchical and network-aware design of Canopus, a proposal message is exchanged between each pair of super-leaves only once. This approach places a lower load on wide-area network links compared to the broadcast-based approach used in EPaxos.

**8.2.1 Writes to Reads Ratio in the Workload.** In this section, we evaluate the performance of Canopus and EPaxos with different ratios of writes in the workload. We run these experiments with nine nodes deployed in three datacenters. Figure 7 shows the results when the workload consists of 1%, 20%, and 50% write requests. For EPaxos, our results show the same results with different write-ratios in the workloads. Therefore, we only show the results for 20% writes workload.

As expected, Canopus has higher throughput for more read-heavy workloads, reaching up to 3.6 million requests per second with 1% writes workload, as compared to the 2.65 millions requests per second with 20% writes workload. However, even with a write intensive workload consisting of 50% writes, Canopus achieves 2.5x higher throughput than EPaxos.

## 9 DISCUSSION

This paper describes the current design of Canopus and evaluates its performance in comparison with other protocols. Additional capabilities may be necessary in order to satisfy the requirements of specific applications.

**Full-rack failures:** In Canopus, nodes belonging to the same super-leaf are placed in the same rack. Therefore, catastrophic failures that cause complete rack failures result in super-leaf failures, which halts the consensus process. Although, top-of-rack switch failures are rare [12] and the fault-tolerant design of datacenters [21, 22, 37, 39] reduce the possibility of rack failures, there can be applications that have more stringent availability requirements. For such applications, we are currently investigating a rack-failure recovery mechanism.

**Experiments at large scale:** In order to limit the cost of our experiments, we only evaluated Canopus with up to 27 nodes across seven datacenters. However, we expect Canopus to scale to a large number of nodes by carefully constructing the LOT hierarchy. The number of nodes in Canopus can be increased either by increasing the size of the super-leaves, or by increasing the number of super-leaves in the system. The LOT should be structured such that the time it takes to complete the operations within a super-leaf is less than the round-trip time between super-leaves. If necessary, the height of the tree can be increased to satisfy this condition.




Figure 5: Throughput and median request completion times of ZKCanopus and ZooKeeper when the system consists of 9 nodes (left) and 27 nodes (right). Note that the x-axis is in log scale.




Figure 6: Throughput and median request completion times of Canopus and EPaxos in the multi-datacenter experiment with 20% writes workload.




Figure 7: Performance of Canopus and EPaxos with different ratios of writes in the workload.

**Byzantine fault tolerance:** Private blockchains use consensus protocols to achieve agreement on the ordering and timing of operations between the participating nodes. In many private blockchains, nodes do not trust each other as they can act maliciously, either intentionally or due to being hacked. A common approach to solve the trust problem is to use Byzantine fault-tolerant

(BFT) consensus algorithms. An alternative approach to establish trust is to create a trusted computing environment [4] using hardware assisted secure enclaves. In this way, the nodes can use non-BFT consensus protocols, such as Canopus, to solve the trust issue. We are in the process of extending Canopus to support BFT.

## 10 CONCLUSION

In this paper, we have presented Canopus, a highly parallel, scalable, and network-aware consensus protocol. Evaluation results using our prototype implementation show that Canopus can perform millions of requests per second with 27 nodes in a single datacenter deployment, which is more than 4x higher than EPaxos. Furthermore, our Canopus-integrated version of ZooKeeper increases the throughput of ZooKeeper by more than 16x for read-heavy workloads.

## ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their valuable feedback. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). This work benefited from the use of the SyN Facility at the University of Waterloo.

## REFERENCES

- [1] 2017. Chain Blockchain Infrastructure. (2017). Retrieved June 19, 2017 from <https://chain.com>
- [2] 2017. Hyperledger Blockchain Technologies. (2017). Retrieved June 19, 2017 from <https://www.hyperledger.org>
- [3] 2017. Kadena Permissioned Blockchain. (2017). Retrieved June 19, 2017 from <http://kadena.io>
- [4] 2017. The Coco Framework. (2017). Retrieved October 16, 2017 from <https://github.com/Azure/coco-framework>
- [5] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. 2008. A Scalable, Commodity Data Center Network Architecture. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication (SIGCOMM '08)*. ACM, 63–74.
- [6] André Allavena, Qiang Wang, Ihab Ilyas, and Srinivasan Keshav. 2006. *LOT: A robust overlay for distributed range query processing*. Technical Report. CS-2006-21, University of Waterloo.
- [7] Jason Baker, Chris Bond, James C Corbett, JJ Furman, Andrey Khorlin, James Larson, Jean-Michel Leon, Yawei Li, Alexander Lloyd, and Vadim Yushprakh. 2011. Megastore: Providing scalable, highly available storage for interactive services. In *Proceedings of the Conference on Innovative Data system Research (CIDR)*. 223–234.
- [8] Martin Biely, Zoran Milosevic, Nuno Santos, and Andre Schiper. 2012. S-Paxos: Offloading the leader for high throughput state machine replication. In *Proceedings of the IEEE Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 111–120.
- [9] Huynh Tu Dang, Marco Canini, Fernando Pedone, and Robert Soulé. 2016. Paxos Made Switch-y. *SIGCOMM Computer Communication Review (CCR)* 46, 2 (May 2016), 18–24.
- [10] Huynh Tu Dang, Daniele Sciascia, Marco Canini, Fernando Pedone, and Robert Soulé. 2015. NetPaxos: Consensus at Network Speed. In *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research (SOSR)*. ACM, 5:1–5:7.
- [11] Nathan Farrington and Alexey Andreyev. 2013. Facebook's Data Center Network Architecture. In *IEEE Optical Interconnects Conference*. 5–7.
- [12] Phillipa Gill, Navendu Jain, and Nachiappan Nagappan. 2011. Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications. In *Proceedings of the ACM SIGCOMM Conference (SIGCOMM '11)*. ACM, 350–361.
- [13] Albert Greenberg, James R Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A Maltz, Parveen Patel, and Sudipta Sengupta. 2011. VL2: a scalable and flexible data center network. *Commun. ACM* 54, 3 (2011), 95–104.
- [14] Patrick Hunt, Mahadev Konar, Flavio P. Junqueira, and Benjamin Reed. 2010. ZooKeeper: Wait-free Coordination for Internet-scale Systems. In *Proceedings of the 2010 USENIX Conference on USENIX Annual Technical Conference (USENIX ATC '10)*. USENIX Association, 9.
- [15] Zsolt István, David Sidler, Gustavo Alonso, and Marko Vukolic. 2016. Consensus in a Box: Inexpensive Coordination in Hardware. In *Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI '16)*. USENIX Association, Santa Clara, CA, 425–438. <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/istvan>
- [16] Flavio P Junqueira, Benjamin C Reed, and Marco Serafini. 2011. Zab: High-performance broadcast for primary-backup systems. In *Proceedings of the IEEE/IFIP 41st International Conference on Dependable Systems Networks (DSN '11)*. IEEE, 245–256.
- [17] Leslie Lamport. 1998. The part-time parliament. *ACM Transactions on Computer Systems (TOCS)* 16, 2 (1998), 133–169.
- [18] Leslie Lamport. 2005. *Generalized consensus and Paxos*. Technical Report. MSR-TR-2005-33, Microsoft Research.
- [19] Leslie Lamport. 2006. Fast paxos. *Springer Distributed Computing* 19, 2 (2006), 79–103.
- [20] Jialin Li, Ellis Michael, Naveen Kr. Sharma, Adriana Szekeres, and Dan R. K. Ports. 2016. Just Say NO to Paxos Overhead: Replacing Consensus with Network Ordering. In *Proceedings of 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, 467–483.
- [21] Vincent Liu, Daniel Halperin, Arvind Krishnamurthy, and Thomas Anderson. 2013. F10: A fault-tolerant engineered network. In *In Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI '13)*. USENIX Association, 399–412.
- [22] Vincent Liu, Danyang Zhuo, Simon Peter, Arvind Krishnamurthy, and Thomas Anderson. 2015. Subways: A Case for Redundant, Inexpensive Data Center Edge Links. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT '15)*. ACM, 27:1–27:13.
- [23] Rafik Makhloufi, Guillaume Doyen, Grégory Bonnet, and Dominique Gaïti. 2014. A survey and performance evaluation of decentralized aggregation schemes for autonomic management. *International Journal of Network Management* 24, 6 (2014), 469–498.
- [24] Yanhua Mao, Flavio Paiva Junqueira, and Keith Marzullo. 2008. Menciuz: building efficient replicated state machines for WANs. In *Proceedings of the 8th USENIX conference on Operating systems design and implementation (OSDI '08)*, Vol. 8. 369–384.
- [25] Parisa Jalili Marandi, Marco Primi, and Fernando Pedone. 2012. Multi-ring Paxos. In *Proceedings of the 2012 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN '12)*. IEEE, 1–12.
- [26] Parisa Jalili Marandi, Marco Primi, Nicolas Schiper, and Fernando Pedone. 2010. Ring Paxos: A high-throughput atomic broadcast protocol. In *Proceedings of the 2010 IEEE/IFIP International Conference on Dependable Systems Networks (DSN '10)*. IEEE, 527–536.
- [27] Iulian Moraru, David G. Andersen, and Michael Kaminsky. 2013. There is More Consensus in Egalitarian Parliaments. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles (SOSP '13)*. ACM, 358–372.
- [28] Iulian Moraru, David G Andersen, and Michael Kaminsky. 2014. EPaxos source code. (2014). Retrieved June 19, 2017 from <https://github.com/efficient/epaxos>
- [29] Iulian Moraru, David G Andersen, and Michael Kaminsky. 2014. Paxos Quorum Leases: Fast Reads Without Sacrificing Writes. In *Proceedings of SoCC*. ACM.
- [30] Iulian Moraru, David G Andersen, and Michael Kaminsky. 2014. Paxos quorum leases: Fast reads without sacrificing writes. In *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 1–13.
- [31] Shuai Mu, Lamont Nelson, Wyatt Lloyd, and Jinyang Li. 2016. Consolidating concurrency control and consensus for commits under conflicts. In *Proceedings of USENIX OSDI*. USENIX Association, 517–532.
- [32] Diego Ongaro and John Ousterhout. 2014. In search of an understandable consensus algorithm. In *Proceedings of USENIX ATC*. USENIX Association, 305–320.
- [33] Marius Poke, Torsten Hoefler, and Colin W. Glass. 2017. AllConcur: Leaderless Concurrent Atomic Broadcast. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing (HPDC '17)*. ACM, 205–218.
- [34] Dan R. K. Ports, Jialin Li, Vincent Liu, Naveen Kr. Sharma, and Arvind Krishnamurthy. 2015. Designing Distributed Systems Using Approximate Synchrony in Data Center Networks. In *Proceedings of 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. USENIX Association, Oakland, CA, 43–57. <https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/ports>
- [35] John Risson and Tim Moors. 2006. Survey of research towards robust peer-to-peer networks: search methods. *Computer Networks* 50, 17 (2006), 3485–3521.
- [36] Sajjad Rizvi, Bernard Wong, and Srinivasan Keshav. 2017. *Canopus: A Scalable and Massively Parallel Consensus Protocol*. Technical Report. CS-2017-08, University of Waterloo.
- [37] Brandon Schlinker, Radhika Niranjana Mysore, Sean Smith, Jeffrey C. Mogul, Amin Vahdat, Minlan Yu, Ethan Katz-Bassett, and Michael Rubin. 2015. Condor: Better Topologies Through Declarative Design. In *Proceedings of the ACM SIGCOMM 2015 Conference*. ACM, 449–463.
- [38] Xuanhua Shi, Haohong Lin, Hai Jin, Bing Bing Zhou, Zuoning Yin, Sheng Di, and Song Wu. 2014. GIRAFFE: A scalable distributed coordination service for large-scale systems. In *Proceedings of IEEE Cluster Computing (CLUSTER)*. IEEE, 38–47.
- [39] Arjun Singh, Joon Ong, Amit Agarwal, Glen Anderson, Ashby Armistead, Roy Bannon, Seb Boving, Gaurav Desai, Bob Felderman, Paulie Germano, and others. 2015. Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network. In *Proceedings of SIGCOMM*. ACM, 183–197.
- [40] Pierre Sutra and Marc Shapiro. 2011. Fast genuine generalized consensus. In *Proceedings of the IEEE Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 255–264.
- [41] Robbert Van Renesse, Kenneth P. Birman, and Werner Vogels. 2003. Astrolabe: A Robust and Scalable Technology for Distributed System Monitoring, Management, and Data Mining. *ACM Transactions on Computer Systems (TOCS)* 21, 2 (May 2003), 164–206.
- [42] Marko Vukolić. 2015. The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication. In *International Workshop on Open Problems in Network Security*. Springer, 112–125.