Fast Flexible Paxos: Relaxing Quorum Intersection for Fast Paxos

Heidi Howard
University of Cambridge
heidi.howard@cl.cam.ac.uk

Aleksey Charapko
University of New Hampshire
aleksey.charapko@unh.edu

Richard Mortier
University of Cambridge
richard.mortier@cl.cam.ac.uk

ABSTRACT

Paxos, the de facto standard approach to solving distributed consensus, operates in two phases, each of which requires an intersecting quorum of nodes. Multi-Paxos reduces this to one phase by electing a leader but this leader is also a performance bottleneck. Fast Paxos bypasses the leader but has stronger quorum intersection requirements.

In this paper we observe that Fast Paxos' intersection requirements can be safely relaxed, reducing to just one additional intersection requirement between phase-1 quorums and any pair of fast round phase-2 quorums. We thus find that the quorums used with Fast Paxos are larger than necessary, allowing alternative quorum systems to obtain new tradeoffs between performance and fault-tolerance.

CCS CONCEPTS

• **Theory of computation** → Distributed algorithms.

KEYWORDS

Distributed computing, Distributed consensus, Paxos

ACM Reference Format:

Heidi Howard, Aleksey Charapko, and Richard Mortier. 2021. Fast Flexible Paxos: Relaxing Quorum Intersection for Fast Paxos. In *International Conference on Distributed Computing and Networking 2021 (ICDCN '21)*, January 5–8, 2021, Nara, Japan. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3427796.3427815>

1 INTRODUCTION

Paxos [18, 19] and its variants [7, 23, 25, 26, 29, 32] provide reliable solutions to the problem of distributed consensus [10]. Thanks to their excellent fault-tolerance properties and proven consistency guarantees, these algorithms often underpin the replicated state machines [31] at the heart of many industrial distributed systems, e.g., Chubby [6], CockroachDB [33], and PaxosStore [36].

Traditionally, the Paxos family of algorithms uses majority quorums, guaranteeing that any two sets containing the majority of nodes intersect, ensuring that previously decided values are not lost. Flexible Paxos [14] relaxes the requirement for intersecting quorums in Paxos, proving that quorum intersection is only required between phases, permitting disjoint quorums to be used within each phase. This result enabled subsequent algorithms to improve

performance by adjusting quorums depending on the phase of the algorithm [1, 4, 8, 9, 28, 35].

Paxos is usually implemented using Multi-Paxos [18, 19], an optimization that elects one node to be a *leader*. This single leader can then achieve distributed consensus in just one phase, but unfortunately also becomes a performance bottleneck.

Seeking to improve performance, a new family of *leaderless* consensus algorithms emerged, starting with Fast Paxos [23], which forms the basis for subsequent algorithms including Generalized Paxos [22] and Egalitarian Paxos [26]. Paxos uses the idea of *rounds* in which at most one value can be *proposed*. Fast Paxos introduced the notion of *fast rounds* where multiple values can be safely proposed in the same round. However, such fast rounds require stronger quorum intersection than classical rounds. Specifically, Paxos only requires that any two quorums intersect whereas Fast Paxos also requires that any quorum intersects with any two fast round quorums. Fast Paxos' quorum intersection requirements can be satisfied by requiring fast round quorums to contain at least three-quarters of nodes. Due to this additional quorum intersection requirement, Fast Paxos and its variants cannot directly benefit from Flexible Paxos.

In this paper, we show that the approach of Flexible Paxos can be safely applied to consensus algorithms that rely on stronger quorum intersection by extending Flexible Paxos to Fast Paxos. The resulting algorithm, which we refer to as *Fast Flexible Paxos*, relaxes the quorum intersection requirements of Fast Paxos. Specifically, Fast Flexible Paxos proves that the only additional quorum intersection requirement is between phase-1 quorums and any pair of fast round phase-2 quorums.

Relaxed quorum intersection in Fast Flexible Paxos permits new performance tradeoffs by manipulation of the quorum systems. For example, reducing the size of fast round quorums reduces the contention in the algorithm. This may further improve overall performance on top of an improvement attained by a smaller quorum alone. For instance, we illustrate that the Fast Flexible Paxos with smaller fast quorums achieves up to 10% better latency than Fast Paxos in low conflict scenarios.

2 BACKGROUND

We begin by recapping how distributed consensus is currently solved by Paxos, Flexible Paxos, and Fast Paxos. The relation between these algorithms is shown in Figure 1.

2.1 Paxos

The Paxos algorithm distinguishes between two roles a node can take: a *proposer* and an *acceptor*. A proposer initiates a decision by executing Paxos using a *round*. Rounds are integers allocated to proposers, and each proposer must propose only one value in each round. A proposal is a pair of a round and a value. The algorithm

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.

ICDCN '21, January 5–8, 2021, Nara, Japan

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8933-4/21/01...\$15.00

<https://doi.org/10.1145/3427796.3427815>

Figure 1: Relationships between distributed consensus algorithms. The arrows denote where one algorithm generalizes over another.

runs in two phases per round, each requiring a quorum of acceptors to proceed.

In phase-1, the proposer learns if a value was decided in any previous round by asking the acceptors to send the last proposal they voted for. Acceptors promise not to vote in any smaller round.

In phase-2, the proposer asks acceptors to vote for a value v . If during phase-1 the proposer learned that a value might become decided, then it must use that value for v . Provided that an acceptor has not promised otherwise, it updates its last proposal voted for and acknowledges that to the proposer. Once the proposer has completed phase-2, it will learn that v is decided.

Paxos requires intersection between any two quorums. If Q denotes the set of quorums then this intersection requirement can be expressed as:

$$\forall Q, Q' \in \mathcal{Q} : Q \cap Q' \neq \{\}$$
 (1)

Consider a simple quorum system based solely on the cardinality of quorums. If q denotes the cardinality of quorums in $\mathcal{Q}$ and n denotes the number of acceptors, then we can express Paxos' quorum intersection requirement (Eq. 1) as:

$$2q > n$$
 (2)

Paxos is often used to decide a sequence of values, where the i^{th} instance of Paxos decides the i^{th} value. *Multi-Paxos* improves the algorithm's performance by pre-executing the phase-1 of every instance by the same stable proposer, known as a *leader*. The leader can then decide each value in just two communication steps, compared to the four communication steps needed by Paxos.

2.2 Flexible Paxos

Flexible Paxos differentiates between the quorums for phase-1, $\mathcal{Q}_1$ and phase-2, $\mathcal{Q}_2$. This approach allows the weakening of Paxos' quorum intersection requirement (Eq. 1) to the following:

$$\forall Q \in \mathcal{Q}_1, \forall Q' \in \mathcal{Q}_2 : Q \cap Q' \neq \{\}$$
 (3)

That is, quorum intersection is not required within each of the two phases. As previously, if q_i denotes the cardinality of quorums in $\mathcal{Q}_i$, then we can express Flexible Paxos' quorum intersection requirement (Eq. 3) as:

$$q_1 + q_2 > n$$
 (4)

In Multi-Paxos, phase-1 is executed rarely compared to phase-2 so applications can decrease the phase-2 quorum to improve performance at the cost of decreased fault-tolerance.

2.3 Fast Paxos

In Multi-Paxos, the leader is a bottleneck and other proposers must first send values to the leader, adding a communication step. Fast Paxos addresses these issues by allowing all proposers to propose values directly to the acceptors. Fast Paxos can thus decide a value in one phase, an optimal solution to distributed consensus [21].

Fast Paxos achieves this by introducing *fast* rounds, where it is safe for multiple values to be proposed in the same round. In a fast round, if the leader is free to propose any value in phase-2 then instead of proposing a specific value it proposes a special *any* value to the acceptors. Proposers can then send proposals directly to the acceptors and each acceptor will vote for the first proposal it receives as if it had been sent by the leader.

Fast Paxos needs stronger quorum intersection for fast rounds to ensure safety and progress in the case of conflicts. If $\mathcal{Q}_f$ denotes fast round quorums and $\mathcal{Q}_c$ denotes classic round quorums then Fast Paxos requires:¹

$$\forall Q, Q' \in \mathcal{Q}_c : Q \cap Q' \neq \{\}$$
 (5)

$$\forall Q, Q' \in \mathcal{Q}_f, \forall Q'' \in \mathcal{Q}_c : Q \cap Q' \cap Q'' \neq \{\}$$
 (6)

$$\forall Q, Q', Q'' \in \mathcal{Q}_f : Q \cap Q' \cap Q'' \neq \{\}$$
 (7)

In other words, Fast Paxos requires that any pair of classic round quorums intersect (Eq. 5), that any pair of fast round quorums intersect with any classic round quorum (Eq. 6) and that any three fast round quorums intersect (Eq. 7).

As before, let q_f and q_c denote the cardinality of quorums in $\mathcal{Q}_f$ and $\mathcal{Q}_c$ respectively. We can now express Fast Paxos' quorum intersection requirements (Eq. 5, Eq. 6 & Eq. 7) as:

$$2q_c > n$$
 (8)

$$q_c + 2q_f > 2n$$
 (9)

$$3q_f > 2n$$
 (10)

Fast Paxos suggests using $q_c = \lfloor \frac{2n}{3} \rfloor + 1$ or $q_c = \lfloor \frac{n}{2} \rfloor + 1$ and $q_f = \lceil \frac{3n}{4} \rceil$ to satisfy these requirements.² The larger quorums required by Fast Paxos have been shown to significantly decrease performance compared to Paxos [15].

3 FAST FLEXIBLE PAXOS

Following the approach of Flexible Paxos, we will differentiate between the quorums used for each phase of Fast Paxos.

Recall that quorum intersection is required between the two phases to ensure that a proposer learns in phase-1 any value which may be decided in phase-2. In phase-2, the proposer needs to pick a value from the highest round it learned during phase-1. In Fast Paxos, there may be multiple such values as acceptors may vote for different values during the phase-2 of fast rounds, requiring a proposer to determine which single value (if any) could be decided in previous rounds. Fast Flexible Paxos achieves this by ensuring that each phase-1 quorum intersects with any pair of fast round phase-2 quorums.

¹The original paper asserts that (a) any two quorums intersect and (b) any two fast quorums and any classic or fast quorum intersect. Since (b) already covers the intersection between two fast quorums as well as a fast and classic quorum, we have reduced (a) to any two classic quorums intersect.

²The former is sometimes written as $n = 3f + 1$ and $q_f = q_c = 2f + 1$ where f is the number of faults which can be tolerated [15].

The quorum intersection requirements are the same regardless of whether phase-1 is for a fast or classic round. Therefore Q_1 denotes the phase-1 quorums (fast or classic) whereas Q_{2c} and Q_{2f} denote the phase-2 quorums for classic and fast rounds respectively. The weakened intersection requirements for Fast Paxos are as follows:

$$\forall Q \in Q_1, \forall Q' \in Q_{2c} : Q \cap Q' \neq \{\} \quad (11)$$

$$\forall Q \in Q_1, \forall Q', Q'' \in Q_{2f} : Q \cap Q' \cap Q'' \neq \{\} \quad (12)$$

In other words, we find that quorum intersection is only required between phase-1 quorums and phase-2 classic round quorums (Eq. 11) and between phase-1 quorums and any pair of phase-2 fast round quorums (Eq. 12). Note that quorum intersection is not required between phase-1 quorums, between phase-2 classic round quorums, or between phase-2 classic round quorums and phase-2 fast round quorums.

If q_1 , q_{2f} and q_{2c} denote the cardinality of quorums in Q_1 , Q_{2f} and Q_{2c} respectively then we can express Fast Flexible Paxos' intersection requirements (Eq. 11 & Eq. 12) as:

$$q_1 + q_{2c} > n \quad (13)$$

$$q_1 + 2q_{2f} > 2n \quad (14)$$

4 CORRECTNESS

Fast Flexible Paxos must ensure that at most one value is decided. We can show this by proving the following two properties:

PROPERTY 1. *At most one value is decided per round.*

PROOF. If the round is classic, then at most one proposer can propose (and therefore decide) a value. If the round is fast, then at most one value will be decided as any two fast phase-2 quorums will intersect (Eq. 12). $\square$

PROPERTY 2. *A proposer will only propose a value in a given round if no smaller round can decide a different value.*

PROOF. Assume that a value v is decided in round r . Consider the next round r' ($r' > r$) where a value is proposed. In phase-1, the proposer of round r' will ask the acceptors to promise not to vote in any smaller rounds and to reply with the last proposal they voted for.

Due to quorum intersection between the phase-1 and phase-2 (Eq. 11 & Eq. 12), at least one acceptor will reply to the proposer in round r' with value v and round r . This is because the acceptor must have voted for value v in round r before participating in round r' as it promises not to vote in any smaller rounds. The acceptor also cannot have voted in any round since r as round r' is the first round after r where a value is proposed. The proposer in round r' will propose the value v as it will not receive any proposals from rounds greater than r . If the round r is fast then the proposer in round r' may receive multiple values from round r , however, the proposer will choose value v since v is decided.

Consider the next round r'' ($r'' > r'$) where a value is proposed. The value proposed in round r'' must be value v as only the value v has been proposed since value v was decided in round r . By induction, we can see that for all rounds larger than r , if a value is proposed then that value will be v . $\square$

Fast Flexible Paxos must also ensure liveness to solve distributed consensus, and in particular, it must satisfy the following property:

PROPERTY 3. *Upon completion of phase-1, a proposer can determine at least one value which is safe to propose in phase-2.*

PROOF. Consider a proposer that has just completed phase-1 of round r' . A value is safe to propose in round r' only if the proposer knows that no smaller round has decided a different value (Property 2). If a proposer receives multiple proposals in phase-1 then it proposes the value with the greatest round r ($r < r'$). However, if the round r is a fast round the proposer may receive multiple values and so must determine which of the values (if any) could be decided in round r . Note that a value could be decided in round r only if there exists a phase-2 fast round quorum of acceptors which may have voted for the value in round r .

For every pair of phase-2 fast round quorums, at least one acceptor which will reply to a proposer in phase-1 of round r' must also vote in both quorums if both quorums decide a value in round r (Eq. 12). The acceptor will only vote for one value in round r and thus will reply to the proposer with only one value. The proposer thus learns that the other value cannot have been decided in round r by any quorum containing that acceptor.

Once the proposer has heard from a phase-1 quorum of acceptors, the proposer can safely eliminate either all or all but one of the values received with round r . $\square$

We adapted the Fast Paxos specification [23] to model check a formal specification of Fast Flexible Paxos using TLA+ [20]. Both our specification and model checking configurations are available online [2].

5 IMPLICATIONS

The weakened intersection requirements show that phase-1 of a fast round can use the same quorum as phase-1 of a classic round. Since the requirement of fast round quorums is stricter than classic round quorums then fast round quorums must be at least as large as classic round quorums.

For example, Fast Paxos suggests using $q_f = \lceil \frac{3n}{4} \rceil$ and $q_c = \lfloor \frac{n}{2} \rfloor + 1$, but our relaxed intersection requirements demonstrate that a simple majority of acceptors is sufficient for phase-1 of fast rounds. Similarly, Fast Paxos also suggests using $q_c = q_f = \lfloor \frac{2n}{3} \rfloor + 1$ and again we observe that is conservative and only one third of acceptors are needed for phase-2 of classic rounds ($q_{2c} = \lceil \frac{n}{3} \rceil$).

More generally, by weakening the intersection requirements of Fast Paxos, we provide more flexibility to choose quorum systems and tradeoffs. In a stable system, phase-1 is rarely executed compared to phase-2 so we can decrease the size of our phase-2 quorums, fast and classic, provided we increase the size of our phase-1 quorums. For example, a system of 11 acceptors could use phase-2 quorums of 7 acceptors for fast rounds and 3 acceptors for classic rounds, if it uses quorums of 9 acceptors for phase-1.

Note that the liveness of such a system does depend upon both phase-1 and phase-2 quorums. For example, we could minimize fast round phase-2 by using a simple majority for q_{2f} , but this would require all acceptors to start a new round.

(a) Instance latency in conflict-free workload at 1400 requests/s.

(b) Instance latency in a workload with 0.5%-1.5% conflicts at 2700 requests/s.

(c) Probability of conflict under different intervals between potentially conflicting commands.

Figure 2: Evaluation of Fast Flexible Paxos in a Paxi 11 node cluster. Fast Paxos uses quorums of the following sizes: $q_c = 6$ and $q_f = 9$. Fast Flexible Paxos runs with $q_1 = 9$, $q_{2f} = 7$, and $q_{2c} = 3$ quorums.

6 PRELIMINARY EVALUATION

The main contribution of this paper is the observation that the quorum intersection requirements of Fast Paxos can safely be relaxed. We have also implemented Fast Flexible Paxos to illustrate

the potential performance improvements this result enables, even with a simple quorum system based solely on quorum cardinality.

We evaluated Fast Flexible Paxos with the aforementioned quorum configuration ($q_1 = 9$, $q_{2f} = 7$, and $q_{2c} = 3$) using Paxi [3] on AWS EC2 m5a.large VMs. We focused on two key aspects of Fast Flexible Paxos: latency and conflict reduction due to the smaller phase-2 fast round quorums. We compared our algorithm against a Fast Paxos ($q_c = 6$ and $q_f = 9$) baseline.

In Figure 2a we illustrate the performance of two algorithms under a workload of 1400 requests/second with no conflicts. Smaller fast round quorums allowed Fast Flexible Paxos to reduce the average and median latency by 5–8% compared to Fast Paxos.

We also evaluated under conflicts by generating a workload with several clients racing to propose different commands for the same consensus instance. For this workload we generated a steady stream of operations with only small intervals between them. We also pre-assigned each operation to an instance to control the potential for conflicts. In about 10% of the cases, we assigned the same instances to two consecutive operations, creating a race condition between them. In this setup there are two possible outcomes for such races: one of the operations reaches the fast round quorum, causing the second to abort; or none of the operations reach the fast round quorum, causing entry into the conflict resolution phase. We then measured the conflict avoidance ratio to study the impact of the Fast Flexible Paxos on conflict handling.

We found that Fast Flexible Paxos entered the conflict recovery almost one-third as frequently as Fast Paxos due to the smaller fast quorum. However, the overall frequency of recovery phases increased substantially for both algorithms as the throughput rises and the interval between the commands shrinks, as Figure 2c shows. Considering the overall performance in the conflict workload, Figure 2b shows that our Fast Flexible Paxos continues to maintain a roughly 5% latency advantage over Fast Paxos even under high load compared to our non-conflict experiment.

We believe Fast Flexible Paxos will enable further performance improvements if quorum systems are used that are not based solely on quorum cardinality [11, 12, 16, 27, 30]. This has already proven to be the case for Flexible Paxos [1, 4, 8, 9, 28, 35]. In particular, Fast Flexible Paxos can benefit from the existing literature on Byzantine and Refined quorum systems [13, 24] as these quorum systems provide stronger quorum intersection.

7 SUMMARY

Fast Paxos allows any proposer to decide a value in two communication steps in the absence of collisions. This is the optimal number of communication steps for distributed consensus. However, to achieve this it needs a stronger quorum intersection than Paxos and thus has not benefited from recent work on relaxing quorum intersection requirements.

Fast Flexible Paxos weakens Fast Paxos' quorum intersection requirements by differentiating between the quorums used in each phase of the algorithm. We find that quorum intersection is only required between any phase-1 quorum and both (a) any phase-2 classic round quorum, and (b) any pair of phase-2 fast round quorums. This shows that the quorum systems used by Fast Paxos

are conservative and that alternative quorum systems could be safely used.

More generally, we have proven that the approach of Flexible Paxos generalizes to distributed consensus algorithms beyond Paxos. We hope more consensus algorithms, particularly those which extend Fast Paxos such as Generalized Paxos [22], Egalitarian Paxos [26], MDCC [17], Alvin [34] and Caesar [5], adopt this approach to relax their quorum intersection requirements, giving applications greater flexibility to determine their performance and fault-tolerance tradeoffs.

ACKNOWLEDGMENTS

This work was funded in part by EPSRC EP/M02315X/1.

REFERENCES

- [1] 2020. LogDevice: Distributed storage for sequential data. <https://logdevice.io>.
- [2] 2020. TLA+ Specification of Fast Flexible Paxos. <https://github.com/fpaxos/fpaxos-tlaplus>.
- [3] Ailidani Ailijiang, Aleksey Charapko, and Murat Demirbas. 2019. Dissecting the Performance of Strongly-Consistent Replication Protocols. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD '19)*. ACM, Amsterdam, Netherlands, 1696–1710. <https://doi.org/10.1145/3299869.3319893>
- [4] Ailidani Ailijiang, Aleksey Charapko, Murat Demirbas, and Tevfik Kosar. 2020. WPaxos: Wide area network flexible consensus. *IEEE Transactions on Parallel and Distributed Systems* 31, 1 (2020), 211–223.
- [5] Balaji Arun, Sebastiano Peluso, Roberto Palmieri, Giuliano Losa, and Binoy Ravindran. 2017. Speeding up Consensus by Chasing Fast Decisions. In *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 49–60.
- [6] Mike Burrows. 2006. The Chubby lock service for loosely-coupled distributed systems. In *7th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*.
- [7] Lásaro Jonas Camargos, Rodrigo Malta Schmidt, and Fernando Pedone. 2007. Multicoordinated Paxos. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC '07)*. ACM, New York, NY, USA, 316–317. <https://doi.org/10.1145/1281100.1281150>
- [8] Michael Eischer, Benedikt Straßner, and Tobias Distler. 2020. Low-Latency Geo-Replicated State Machines with Guaranteed Writes. In *Proceedings of the 7th Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC '20)*. ACM, Heraklion, Greece, Article 13, 9 pages. <https://doi.org/10.1145/3380787.3393686>
- [9] Vitor Enes, Carlos Baquero, Tuanir França Rezende, Alexey Gotsman, Matthieu Perrin, and Pierre Sutra. 2020. State-Machine Replication for Planet-Scale Systems. In *Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys '20)*. ACM, Heraklion, Greece, Article 24, 15 pages. <https://doi.org/10.1145/3342195.3387543>
- [10] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. 1985. Impossibility of Distributed Consensus with One Faulty Process. *J. ACM* 32, 2 (April 1985), 374–382. <https://doi.org/10.1145/3149.214121>
- [11] Hector Garcia-Molina and Daniel Barbara. 1985. How to Assign Votes in a Distributed System. *J. ACM* 32, 4 (Oct. 1985), 841–860. <https://doi.org/10.1145/4221.4223>
- [12] David K. Gifford. 1979. Weighted Voting for Replicated Data. In *Proceedings of the Seventh ACM Symposium on Operating Systems Principles (SOSP '79)*. ACM, Pacific Grove, CA, USA, 150–162. <https://doi.org/10.1145/800215.806583>
- [13] Rachid Guerraoui and Marko Vukolić. 2007. Refined Quorum Systems. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC '07)*. ACM, Portland, OR, USA, 119–128. <https://doi.org/10.1145/1281100.1281120>
- [14] Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. 2017. Flexible Paxos: Quorum Intersection Revisited. In *20th International Conference on Principles of Distributed Systems (OPODIS 2016)*. 25:1–25:14. <https://doi.org/10.4230/LIPIcs.OPODIS.2016.25>
- [15] Flavio Junqueira, Yanhua Mao, and Keith Marzullo. 2007. Classic Paxos vs. Fast Paxos: Caveat Emptor. In *Proceedings of the 3rd Workshop on on Hot Topics in System Dependability (HotDep'07)*. USENIX Association, Edinburgh, UK, 18–es.
- [16] Flavio P. Junqueira and Keith Marzullo. 2005. The Virtue of Dependent Failures in Multi-Site Systems. In *Proceedings of the First Conference on Hot Topics in System Dependability (HotDep'05)*. USENIX Association, Yokohama, Japan, 1.
- [17] Tim Kraska, Gene Pang, Michael J. Franklin, Samuel Madden, and Alan Fekete. 2013. MDCC: Multi-Data Center Consistency. In *Proceedings of the 8th ACM European Conference on Computer Systems (EuroSys '13)*. ACM, Prague, Czech Republic, 113–126. <https://doi.org/10.1145/2465351.2465363>
- [18] Leslie Lamport. 1998. The Part-time Parliament. *ACM Trans. Comput. Syst.* 16, 2 (May 1998), 133–169. <https://doi.org/10.1145/279227.279229>
- [19] Leslie Lamport. 2001. Paxos made simple. *ACM SIGACT News (Distributed Computing Column)* (2001).
- [20] Leslie Lamport. 2002. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [21] Leslie Lamport. 2003. Lower Bounds for Asynchronous Consensus. In *Future Directions in Distributed Computing: Research and Position Papers*. Springer-Verlag, 22–23.
- [22] Leslie Lamport. 2004. *Generalized Consensus and Paxos*. Technical Report MSR-TR-2005-33. Microsoft Research.
- [23] Leslie Lamport. 2005. *Fast Paxos*. Technical Report MSR-TR-2005-112. Microsoft Research.
- [24] Dahlia Malkhi and Michael Reiter. 1998. Byzantine Quorum Systems. *Distrib. Comput.* 11, 4 (Oct. 1998), 203–213. <https://doi.org/10.1007/s004460050050>
- [25] Yanhua Mao, Flavio P. Junqueira, and Keith Marzullo. 2008. Mencius: Building Efficient Replicated State Machines for WANs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation (OSDI'08)*. USENIX Association, San Diego, CA, USA, 369–384.
- [26] Iulian Moraru, David G. Andersen, and Michael Kaminsky. 2013. There is More Consensus in Egalitarian Parliaments. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles (SOSP '13)*. ACM, Farmington, PA, USA, 358–372. <https://doi.org/10.1145/2517349.2517350>
- [27] Moni Naor and Avishai Wool. 1998. The Load, Capacity, and Availability of Quorum Systems. *SIAM J. Comput.* 27, 2 (April 1998), 423–447. <https://doi.org/10.1137/S0097539795281232>
- [28] Faisal Nawab, Divyakant Agrawal, and Amr El Abbadi. 2018. DPaxos: Managing Data Closer to Users for Low-Latency and Mobile Applications. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*. ACM, Houston, TX, USA, 1221–1236. <https://doi.org/10.1145/3183713.3196928>
- [29] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference (USENIX ATC'14)*. USENIX Association, Philadelphia, PA, USA, 305–320.
- [30] David Peleg and Avishai Wool. 1995. Crumbling Walls: A Class of Practical and Efficient Quorum Systems. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing (PODC '95)*. ACM, Ottawa, ON, Canada, 120–129. <https://doi.org/10.1145/224964.224978>
- [31] Fred B. Schneider. 1990. Implementing Fault-Tolerant Services Using the State Machine Approach: A Tutorial. *ACM Comput. Surv.* 22, 4 (Dec. 1990), 299–319. <https://doi.org/10.1145/98163.98167>
- [32] Pierre Sutra and Marc Shapiro. 2011. Fast Genuine Generalized Consensus. In *Proceedings of the 2011 IEEE 30th International Symposium on Reliable Distributed Systems (SRDS '11)*. IEEE Computer Society, 255–264. <https://doi.org/10.1109/SRDS.2011.38>
- [33] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis, Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, Paul Bardea, Amruta Ranade, Ben Darnell, Bram Gruneir, Justin Jaffray, Lucy Zhang, and Peter Mattis. 2020. CockroachDB: The Resilient Geo-Distributed SQL Database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD '20)*. ACM, Portland, OR, USA, 1493–1509. <https://doi.org/10.1145/3318464.3386134>
- [34] Alexandru Turcu, Sebastiano Peluso, Roberto Palmieri, and Binoy Ravindran. 2014. Be General and Don't Give Up Consistency in Geo-Replicated Transactional Systems. In *Principles of Distributed Systems*. Springer International Publishing, Cham, 33–48.
- [35] Muhammed Uluyol, Anthony Huang, Ayush Goel, Mosharaf Chowdhury, and Harsha V. Madhyastha. 2020. Near-Optimal Latency Versus Cost Tradeoffs in Geo-Distributed Storage. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. USENIX Association, Santa Clara, CA, 157–180. <https://www.usenix.org/conference/nsdi20/presentation/uluyol>
- [36] Jianjun Zheng, Qian Lin, Jiatao Xu, Cheng Wei, Chuwei Zeng, Pingan Yang, and Yunfan Zhang. 2017. PaxosStore: high-availability storage made practical in WeChat. *Proceedings of the VLDB Endowment* 10, 12 (2017), 1730–1741.