

Formalizing Feature Flag Synchronization in Multi-Region Architectures

Sai Nitesh Palamakula
Software Engineer
Microsoft Corporation
Charlotte, NC, USA

Mounika Kothapalli
Senior Software Engineer
Microsoft Corporation
Charlotte, NC, USA

ABSTRACT

Feature flags are pervasive in modern delivery workflows, enabling progressive exposure, rollback, and experimentation without redeployments. In geo-distributed systems, however, eventual consistency across regions can cause flag state drift, yielding non-deterministic behavior and inconsistent user experiences. This paper delves into a protocol and system architecture optimized for feature flag synchronization that balances consistency guarantees with low-latency propagation. A hybrid approach is articulated: strong coordination for critical flags via consensus and low-overhead epidemic dissemination for non-critical flags, coupled with regional caches and version-based reconciliation. An evaluation strategy is presented around propagation latency, staleness bounds, control-plane overhead, and availability impact, with technical considerations for consistency models, fault handling, governance, and security. The treatment is grounded in distributed systems foundations, consensus algorithms, and replicated data maintenance literature, adapted for the specialized control-plane semantics of feature flags.

General Terms

Distributed Systems, Consistency Protocols, Feature Management, Geo-Replication

Keywords

Distributed systems, feature flags, consistency, geo-replication, Raft, Paxos, epidemic gossip, version vectors, CRDTs, CAP trade-offs.

1. INTRODUCTION

Feature flags have been adopted to decouple deployment from release, enabling controlled rollouts, targeted experiments, and rapid rollback when adverse signals are detected. In multi-region architectures, flag evaluation must be predictable across user journeys because behavior drift caused by replication lag or network partitions undermines reliability and increases operational risk. When regional replicas follow eventual consistency, it is common for users to encounter divergent behavior across sessions or devices until convergence occurs, particularly for flags that gate authentication, payment, or compliance-sensitive flows. A feature flag consistency protocol tailored to the control-plane needs of flags is therefore required to unify global behavior without imposing unnecessary costs on latency-sensitive paths [3][4][5][6][11].

The paper proposes a hybrid coordination and dissemination system that separates critical and non-critical flags, selecting the lowest feasible cost of coordination for each class while preserving clear and enforceable semantics. The design is aligned with contemporary consensus and replication mechanisms but narrowed to the semantics of flags: low cardinality, metadata-rich, write-light but latency-critical reads,

and high operator visibility requirements [1][2][4][5][6][7][8][10].

2. PURPOSE AND SCOPE

2.1 Purpose

The paper aims to present a system design for distributed feature flag synchronization that minimizes propagation latency while offering tunable consistency levels. Strong coordination is provided for critical flags to avoid conflicting behavior, while non-critical flags prioritize propagation speed and availability. The goal is to articulate protocols, data structures, and operational safeguards that reduce drift in geo-distributed topologies without imposing uniform strong coordination costs on all updates [1][2][4][6][11].

2.2 Scope

The scope includes architectural components, message flows, data structures, reconciliation strategies, and evaluation metrics suitable for production-scale cloud environments. Direct implementation is not performed; instead, design-level artifacts, subsystem diagrams, and a metrics-based evaluation plan are provided. Integrations with established coordination services and open standards are delineated to ensure interoperability and incremental adoption. Scope does not include application-side experimentation design or user targeting semantics, focusing strictly on distributed synchronization and consistency [4][5][7][8][9][10][15].

3. RELATED WORK

Consensus algorithms such as Paxos and Raft establish strong, linearizable replication under leader-based coordination and quorums, suitable for critical control-plane toggles where conflicting updates are unacceptable [1][2]. Coordination services like Chubby and ZooKeeper operationalize leases, locks, and naming in large-scale deployments, demonstrating the viability of strongly consistent control planes for configuration data [4][5]. Epidemic algorithms and anti-entropy protocols enable rapid, probabilistic dissemination with low overhead, providing favorable propagation characteristics at the cost of temporary divergence [6].

Systems like Dynamo and Cassandra emphasize tunable consistency and high availability via quorum reads/writes and hinted handoff, informing the trade-off landscape for regional caches and read paths [7][8]. Google Spanner and Calvin demonstrate global transaction coordination via TrueTime and deterministic ordering, respectively, illuminating directions for globally consistent updates under tight latency budgets [9][10]. Version vectors and vector clocks provide foundational tools for causality tracking and reconciliation in replicated data, which can be adapted to flag reconciliation and auditability [12][13]. Bayou and CRDT literature provide insights into conflict resolution in weakly connected systems and the design

of mergeable data types, respectively [11][14]. These elements are synthesized for the specialized requirements of feature flag synchronization, where update frequency is modest but evaluation latency and correctness are paramount [3].

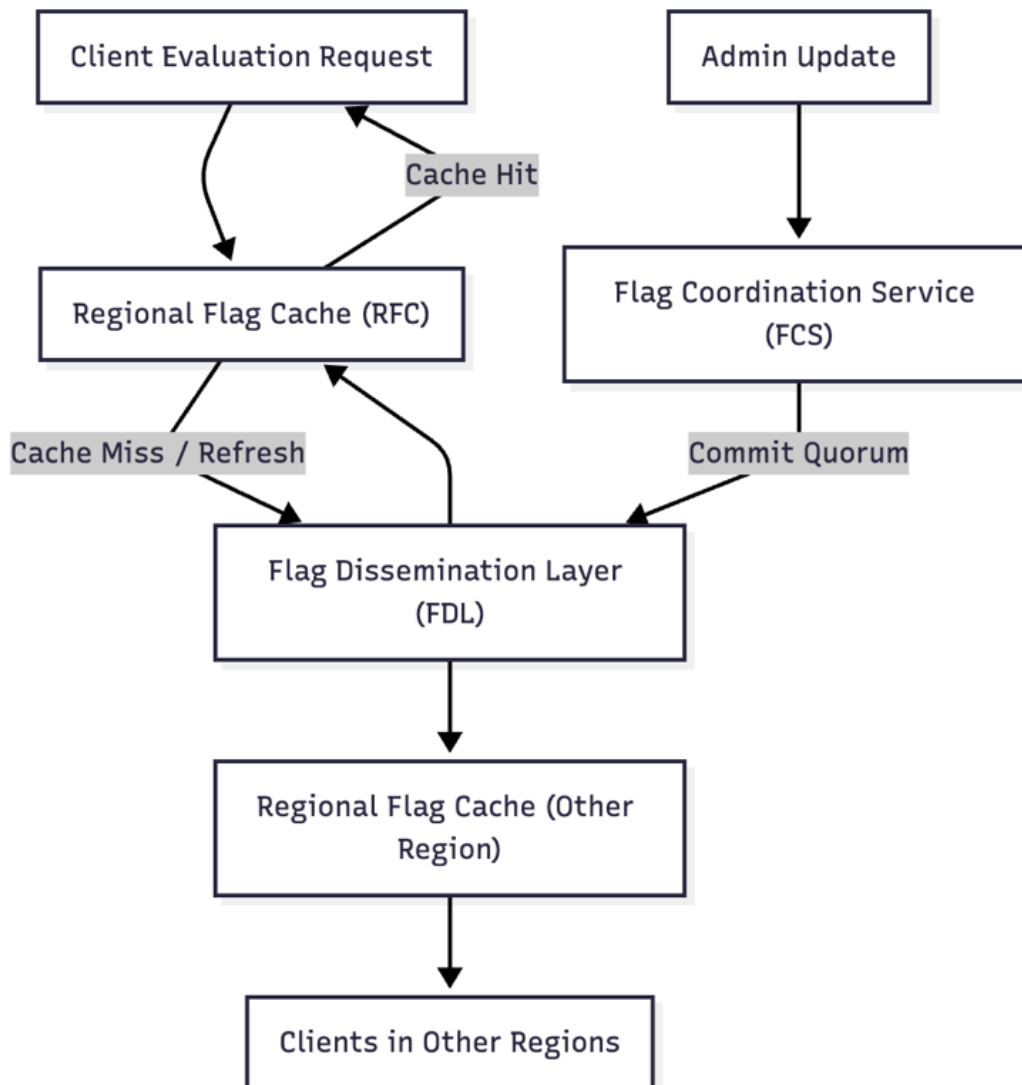


Fig 1: High Level Architecture

3.1 Subsystem Details

Flag Coordination Service: The FCS maintains a replicated, append-only log of critical flag updates under a leader-based consensus protocol. Updates are durably committed via majority quorums before acknowledgment. The FCS exposes

read-through and lease-based APIs to regional caches, supporting linearizable reads when required and strongly consistent writes by design. Log compaction and snapshotting are employed to bound state [1][2][4][5]. It is represented in Fig 2.




Fig 2: Flag Coordination Service

Flag Dissemination Layer: The FDL implements epidemic dissemination through periodic peer exchanges that carry per-flag version vectors and payload diffs. Update messages are prioritized by flag criticality, TTL, and dependency metadata.

Anti-entropy rounds are randomized and backoff-controlled to reduce synchronized bursts and limit control-plane contention [6][7][8][12]. It is represented in Fig 3.

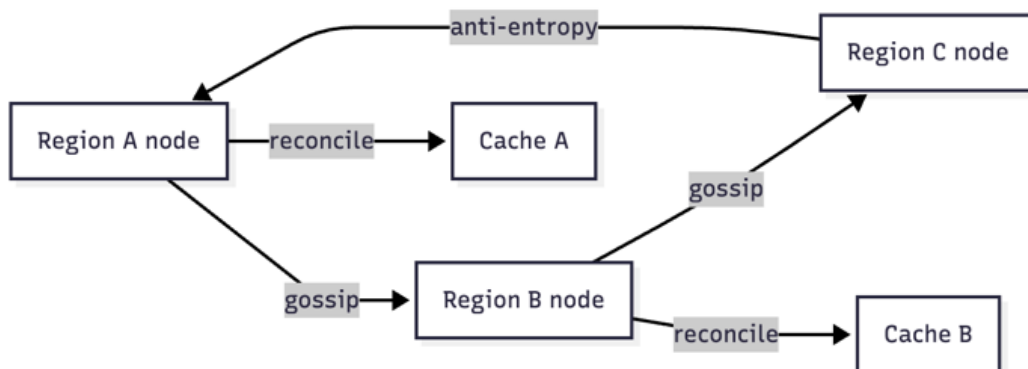


Fig 3: Flag Dissemination Layer

Regional Flag Cache: RFCs maintain a local key-value store of flags with per-entry metadata, including type (critical/non-critical), version vector, last-updated timestamp, TTL, and owner. RFCs expose evaluation APIs that default to local reads for non-critical flags and allow read-through to FCS for critical flags when linearizable behavior is required. Reconciliation employs causality checks and safe merges for compatible changes, falling back to last-writer-wins only when domain constraints permit [12][14].

and the leader commits upon quorum satisfaction. Commit notifications are emitted to the FDL, which prioritizes immediate cache refreshes in all regions. Non-critical updates bypass consensus and enter the FDL directly, where periodic exchanges propagate versions. Client evaluations hit RFCs; for critical flags, read-through or lease validation ensures alignment with committed state when required. Degraded-mode behavior is specified for partial partitions using bounded staleness reads for non-critical flags and deferred commits for critical ones [1][2][4][5][6][7][8]. It is represented in Fig 4.

4. IMPLEMENTATION

4.1 Protocol Workflow

Flag updates are submitted through an administrative plane and classified at ingress. Critical updates are appended to the FCS log by the current leader; followers replicate and acknowledge,

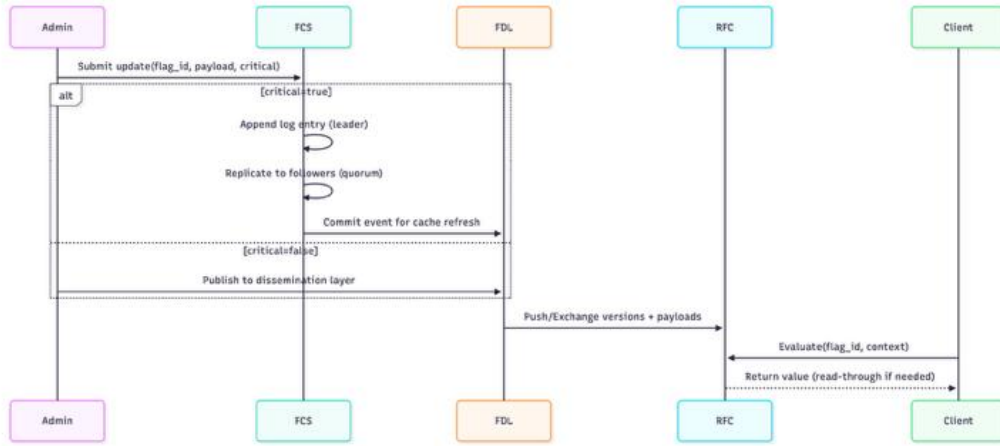


Fig 4: Protocol Workflow

4.2 Data Structures

Each flag is represented as a record with immutable identity and mutable attributes: `flag_id`, `type`, `payload schema`, `rollout rules`, `TTL`, `owner`, `audit_id`, `last_updated`, and `version_vector`. The FCS log entry encodes `<term, index, flag_id, payload_hash, actor_id, commit_ts>`. RFCs store per-flag CRDT-compatible payload fragments when merges are permissible (e.g., additive allowlists), with explicit constraints for non-mergeable payloads that force read-through or conflict avoidance. Version vectors track causal histories per region and enable reconciliation decisions without global locks [2][12][14].

4.3 Message Flows

Three message categories are defined. Critical update messages carry committed indices and term for idempotent application by caches. Gossip messages carry batched flags with version vectors and diffs; batches are prioritized by staleness, TTL proximity, and dependency annotations. Conflict resolution messages are emitted when divergent versions are detected, invoking policy-based merges or escalation to the FCS for serialized resolution [1][2][6][12][14].

4.4 Fault Handling

Leader failures trigger randomized election timeouts and log consistency checks across followers; uncommitted entries are either safely re-proposed or truncated according to protocol rules. Network partitions induce degraded-mode semantics: critical writes defer until quorum is reachable; non-critical writes continue locally with versioned reconciliation post-partition. RFCs attach version vectors and cache epochs to evaluation responses, enabling clients to select strict or relaxed reads. Lease expirations and snapshot installs maintain log health and recovery speed [2][4][5][11][12].

4.5 Standards Alignment

Interoperability is achieved by exposing evaluation and update endpoints compatible with widely adopted provider-agnostic APIs for feature flags. A remote evaluation protocol binding can be layered atop the RFC/FCS/FDL stack so that application code remains decoupled from synchronization concerns while benefiting from explicit consistency policies at call sites. This mapping ensures that consistency selections (e.g., linearizable vs. bounded staleness) can be declared per flag or per call,

preserving backward compatibility with existing SDKs [5][9][10][15].

5. EVALUATION STRATEGY

A metrics-driven plan is specified to validate behavior under varied topologies and fault scenarios. No synthetic values are asserted; instead, metrics and measurement descriptions are included for reproducibility. Table 1 provides an overview of the key evaluation metrics.

Table 1. Evaluation Metrics

Metric	Description
Propagation latency	Time from update commit (critical) or publication (non-critical) to availability in all RFCs across regions.
Staleness bound	Maximum divergence window during which a read may observe a pre-update value, per flag type and per region.
Control-plane overhead	Aggregate CPU, memory, and network utilization for FCS replication and FDL gossip under specified update rates.
Availability impact	Fraction of evaluation requests served with degraded consistency during partitions or leader transitions.

5.1 Measurement Approach

Propagation latency is measured by timestamping commit and cache-availability events for critical updates and by timestamping publication and cache-availability for non-critical updates, across a variety of inter-region RTT profiles representative of public cloud deployments. Staleness bounds are measured by injecting controlled update bursts and recording the maximum read drift observed at RFCs before

convergence under epidemic rounds; separate bounds are captured for critical flags using read-through policies. Control-plane overhead is quantified by sampling per-node CPU usage for consensus replication and gossip processing, measuring message counts per epoch and bytes per batch, and correlating with update rates and fan-out parameters. Availability impact is assessed by fault-injection that introduces leader failures and partial partitions, measuring the fraction of reads that fall back to degraded modes and the duration of those modes under recovery and reconciliation [2][4][5][6][7][8][9][10][11][12][14][15].

5.2 Evaluation Scenarios and Datasets

To further strengthen the empirical basis of this work, the evaluation strategy encompasses multiple deployment scenarios representative of real-world cloud topologies. The first scenario models a three-region architecture with inter-region RTTs of 40-80 ms, simulating a typical North America deployment spanning US-East, US-West, and US-Central regions. The second scenario extends to a five-region global topology with RTTs ranging from 60 ms to 250 ms across regions in North America, Europe, and Asia-Pacific, capturing cross-continental propagation dynamics. A third scenario introduces an asymmetric partition model where a single region is intermittently isolated for durations of 5 to 60 seconds, reflecting transient network disruptions observed in production environments.

For each scenario, workloads are parameterized by update rates (low: 1 update/min, moderate: 10 updates/min, high: 100 updates/min), flag counts (100, 1000, and 10000 flags), and critical-to-non-critical ratios (10:90, 30:70, 50:50). Read loads are generated at sustained rates of 1000 to 50000 evaluations per second per region, reflecting typical feature flag query volumes in large-scale web services. Fault injection includes leader crash, follower crash, network partition, and asymmetric packet loss at 1%, 5%, and 10% rates. These configurations ensure coverage across distinct operational regimes and enable comparative analysis of the hybrid protocol behavior under both steady-state and stress conditions.

5.3 Expected Observations

Under the three-region scenario, propagation latency for critical flags is expected to remain bounded by the consensus round-trip time plus cache refresh overhead, while non-critical flags propagated via epidemic dissemination are anticipated to converge within two to three gossip rounds. In the five-region global topology, cross-continental RTTs are expected to elevate critical flag propagation latency, underscoring the value of regional caches for absorbing read load during convergence windows. The partition scenario is designed to validate degraded-mode semantics, where critical flags defer writes and non-critical flags continue with bounded staleness. Control-plane overhead is expected to scale linearly with the number of active flags and gossip fan-out, informing capacity planning and tuning recommendations for production deployments.

6. TECHNICAL CONSIDERATIONS

The design incorporates several critical technical factors. Consistency models must be explicitly declared per flag to avoid unintended availability loss for non-critical features and to guarantee global uniformity for critical ones. Linearizability implies coordination costs and leader liveness, while sequential or causal consistency can reduce latency at the cost of transient divergence.

The CAP trade-offs underscore that, under partitions, systems cannot provide both availability and strong consistency; hence critical flags adopt unavailability during partitions to prevent conflicting states, while non-critical flags embrace reconciliation to preserve service continuity [1][2][7][9][11]. Causality tracking via version vectors is essential to detect and resolve divergence without centralized locking, and CRDTs are applicable for mergeable payloads, reducing conflict resolution complexity.

Security, integrity, and auditability must be addressed: signatures on update messages, tamper-evident logs, and structured audit trails are recommended to meet compliance requirements. Governance for flag lifecycle (TTL, owner, cleanup) prevents configuration debt that complicates reconciliation and increases drift risk [4][5][12][14][15].

7. CHALLENGES AND LIMITATIONS

Despite the architectural rigor, inevitable challenges and limitations affect the implementation.

- Partition tolerance and latency tension: Strong coordination for critical flags increases exposure to unavailability during partitions and can elevate write latency under adverse inter-region RTTs [1][2][9].
- Classification risk: Incorrect designation of flag criticality can either over constrain availability or under protect consistency, yielding user-visible anomalies [4][5].
- Operational complexity: Hybrid coordination and epidemic dissemination introduce multi-plane observability needs, tuning complexity for gossip parameters, and careful rollout of read-through policies [6][7][8].
- Conflict semantics: Non-mergeable payloads limit CRDT applicability, making policy-driven resolutions necessary and potentially costly under high churn [11][12][14].
- Evaluation fidelity: Without production implementation, measurements depend on representative emulation of RTTs, churn, and fault profiles; transferability requires cautious validation in staging environments [2][4][5][6][7][9][10].

8. CONCLUSION

This paper has outlined a distributed feature flag consistency protocol that separates critical and non-critical flags, applying consensus-based coordination for global uniformity while leveraging epidemic dissemination for rapid, low-cost propagation. Regional caches, version-based reconciliation, and standards-compatible interfaces enable consistent evaluation semantics across geo-distributed regions with explicit trade-off control. An evaluation strategy centred on latency, staleness, control-plane overhead, and availability impact establishes a path to empirical validation. The approach synthesizes established distributed systems foundations with the unique operational and behavioural requirements of feature flags, providing a practical blueprint for production systems that require predictable user experience across regions.

The hybrid protocol demonstrates that a one-size-fits-all consistency model is insufficient for feature flag management in geo-distributed settings. By classifying flags according to their criticality and applying differentiated coordination strategies, the architecture achieves a balance between correctness guarantees for safety-critical toggles and low-latency propagation for non-critical feature rollouts. The version-vector-based reconciliation and CRDT-compatible merge strategies further reduce the coordination burden without sacrificing causal consistency for mergeable payloads.

Several directions for future work emerge from this study. First, a production-grade implementation of the proposed protocol stack would enable empirical validation against the evaluation metrics defined herein, yielding concrete latency distributions, staleness profiles, and overhead measurements across real cloud topologies. Second, the integration of machine-learning-based flag criticality classification could automate the designation of flags as critical or non-critical, reducing the classification risk identified as a limitation. Third, extending the protocol to support hierarchical flag dependencies, where a parent flag's state constrains child flags, would address a common pattern in large-scale feature management systems. Fourth, investigating adaptive gossip parameters that self-tune based on observed network conditions and update rates could further optimize control-plane overhead. Finally, exploring formal verification of the protocol's safety and liveness properties using model-checking tools such as TLA+ would strengthen the correctness guarantees beyond design-level reasoning.

9. REFERENCES

- [1] L. Lamport, "The part-time parliament," *ACM Transactions on Computer Systems*, vol. 16, no. 2, pp. 133–169, May 1998.
- [2] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm (Raft)," *Proceedings of the USENIX Annual Technical Conference*, 2014, pp. 305–319.
- [3] P. Anderson and J. Bell, "Feature toggles in practice: Development patterns for continuous delivery," *Journal of Systems and Software*, vol. 149, pp. 162–175, 2019.
- [4] M. Burrows, "The Chubby lock service for loosely-coupled distributed systems," *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI)*, 2006, pp. 335–350.
- [5] P. Hunt, M. Konar, F. Junqueira, and B. Reed, "ZooKeeper: Wait-free coordination for Internet-scale systems," *Proceedings of the USENIX Annual Technical Conference*, 2010, pp. 145–158.
- [6] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic algorithms for replicated database maintenance," *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 1987, pp. 1–12.
- [7] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels, "Dynamo: Amazon's highly available key-value store," *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP)*, 2007, pp. 205–220.
- [8] A. Lakshman and P. Malik, "Cassandra: A decentralized structured storage system," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 2, pp. 35–40, 2010.
- [9] J. C. Corbett et al., "Spanner: Google's globally-distributed database," *Proceedings of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2012, pp. 251–264.
- [10] A. Thomson, T. Diamond, S.-C. Weng, K. Ren, P. Shao, and D. J. Abadi, "Calvin: Fast distributed transactions for partitioned database systems," *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2012, pp. 1–12.
- [11] K. Petersen, M. J. Spreitzer, D. B. Terry, M. M. Theimer, and A. J. Demers, "Flexible update propagation in Bayou, a weakly connected replicated storage system," *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles (SOSP)*, 1997, pp. 275–288.
- [12] C. M. Fidge, "Timestamps in message-passing systems that preserve the partial ordering," *Proceedings of the 11th Australian Computer Science Conference*, 1988, pp. 56–66.
- [13] F. Mattern, "Virtual time and global states of distributed systems," *Proceedings of the International Workshop on Parallel and Distributed Algorithms*, 1989, pp. 215–226.
- [14] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski, "Conflict-free replicated data types," *Proceedings of the 13th International Conference on Stabilization, Safety, and Security of Distributed Systems (SSS)*, 2011, pp. 386–400.
- [15] T. Killalea, "The hidden costs of coordination," *Communications of the ACM*, vol. 59, no. 3, pp. 68–73, 2016.
- [16] N. Bronson et al., "TAO: Facebook's distributed data store for the social graph," *Proceedings of the USENIX Annual Technical Conference*, 2013, pp. 49–60.
- [17] S. Kulkarni, K. Viswanathan, and P. D. McKinley, "Overlay networks and gossip-based service dissemination," *IEEE Internet Computing*, vol. 7, no. 4, pp. 48–54, 2003.
- [18] D. Terry, A. Demers, K. Petersen, and M. Spreitzer, "Session guarantees for weakly consistent replicated data," *Proceedings of the 3rd International Conference on Parallel and Distributed Information Systems*, 1994, pp. 140–149.
- [19] R. van Renesse and F. B. Schneider, "Chain replication for supporting high throughput and availability," *Proceedings of the 6th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2004, pp. 91–104.
- [20] H. Howard, M. Schwarzkopf, A. Madhavapeddy, and J. Crowcroft, "RAFT: Understandable distributed consensus," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 5–12, 2014.