



 Research Article

## Reliable and Scalable Mobile Technology Architecture: A Formal Model, Synchronization Protocol, and Reference Design

Submission Date: June 11, 2021, Accepted Date: July 13, 2021,

Published Date: August 31, 2021 |

Journal Website:

<https://theamericanjournals.com/index.php/tajet>

Srikanth Puram  
Novi, Michigan, USA

Copyright: Original content from this work may be used under the terms of the creative common's attributes 4.0 license.

### ABSTRACT

Mobile systems are distributed systems whose nodes are intermittently connected, frequently restarted by the host OS, energy-constrained, and deployed on hardware the developer does not control. Despite this, much of the architectural discussion around mobile software still treats the network as reliable and the device as a stateless view onto a backend. This paper takes the opposite starting point. We model the mobile client as a first-class replica with its own durable state, and we treat the link to the backend as an unreliable, asynchronous channel that is unavailable for unbounded periods.

From that model we derive two properties of interest — reliability (correct behavior under failure, disconnection, and partial outage) and scalability (growth in users, data, device heterogeneity, and team size without proportional growth in cost or operational risk) — and we show that, contrary to the usual framing, the two goals largely reinforce each other when the architecture is built local-first. The contributions are: (1) a formal system and failure model for mobile distributed systems; (2) a precise treatment of the consistency guarantees achievable at the mobile tier and the conditions under which each holds; (3) a delta-based synchronization protocol with explicit idempotency, conflict-detection, and convergence semantics, presented as algorithms rather than prose; (4) a quantitative treatment of retry, backoff, and backpressure, including thundering-herd mitigation; (5) a layered reference architecture and a security threat model for local-first operation; and (6) a falsifiable evaluation framework, comprising metrics and a fault-injection methodology, illustrated through a worked deployment scenario spanning phones, tablets, and wearables.

## KEYWORDS

mobile architecture, distributed systems, eventual consistency, version vectors, CRDTs, idempotency, offline-first, cross-device sync, fault injection, edge computing

## 1. INTRODUCTION

A mobile application past its prototype stage is not a UI talking to an API. It is a replica in a distributed system. It holds state, mutates that state while disconnected, and must later reconcile with one or more authoritative replicas over a channel that is slow, lossy, and frequently absent. Every hard problem in the field duplicate writes, data that silently diverges between devices, work lost when the OS kills a process mid-operation, releases that brick a fraction of a fleet is a distributed-systems problem wearing mobile clothing.

This paper argues that the right way to get reliability and scalability is to design from the distributed-systems model outward, rather than bolting failure handling onto a request/response client after the fact. We define the two target properties precisely:

- **Reliability.** Given the failure model of Section 2, the system (i) never loses or corrupts a durably acknowledged user write, (ii) converges all replicas to an agreed state once connectivity is restored, and (iii) continues to serve core read and write functionality during disconnection. These are testable conditions, not adjectives.
- **Scalability.** As the number of clients  $N$ , the per-client dataset size  $D$ , the codebase size, and the team size grow, the marginal backend cost per active client, the sync

payload per operation, and the build/merge overhead per engineer should stay bounded ideally flat or sublinear.

The central claim is that a local-first architecture [7], in which the device holds an authoritative replica and the network is an optimization rather than a precondition, satisfies both properties simultaneously, because the same mechanisms that tolerate disconnection (a durable local store, a change log, delta synchronization) also reduce backend load and request volume.

### 1.1 Related Work

The building blocks of this architecture are individually well established. Ordering and consistency rest on Lamport's happens-before relation and logical clocks [1], the availability/consistency trade-off formalized as the CAP theorem [4, 5], the per-session guarantees of Terry et al. [3], and the eventual-consistency model surveyed by Vogels [10]. Replication and convergence draw on version vectors for detecting concurrent updates [2], Conflict-Free Replicated Data Types and their Strong Eventual Consistency conditions [6], and the practical highly-available key-value store design of Dynamo [11], which applied vector clocks at scale. The local-first stance treating the device, rather than the server, as the authoritative replica was articulated by Kleppmann et al. [7]. Structural and operational patterns come from

clean/hexagonal architecture [12], the resilience patterns catalogued by Nygard [9], and the backoff- with-jitter analysis popularized by Brooker [8].

What is largely absent from this prior work is a treatment that composes these results specifically for the mobile tier. The replicated-data and consistency literature is framed around data-center replication or desktop collaboration; the resilience and architecture patterns are framed around server-side services. This paper's contribution is

integrative: an explicit mobile failure model (Section 2), a synchronization protocol that combines version-vector conflict detection with per-type resolution (Section 5), a convergence argument for a schema mixing CRDT, last-writer-wins, and server-arbitrated data types (Section 5.4), and a single reference architecture (Section 7) in which the reliability and scalability mechanisms coincide.

Section 2 gives the system and failure model. Section 3 fixes the consistency vocabulary. Section 4 treats reliability mechanisms with their failure analyses. Section

5 presents the synchronization protocol. Section 6 treats scalability and the quantitative behavior of retries and backpressure. Section 7 is the reference architecture. Section 8 is the security threat model. Section 9 is the evaluation framework and worked example. Section 10 covers open problems.

## 2. SYSTEM AND FAILURE MODEL

We make the model explicit so that every later claim can be checked against it.

### 2.1 Topology

Let the system be a set of replicas  $R = \{r_0, r_1, \dots, r_n\}$ . One distinguished replica  $r_0$  is the **server replica** (logically; it may itself be a horizontally scaled, internally replicated service). The remaining  $r_i$  are **client replicas** running on devices. Each  $r_i$  holds a local state  $S_i$ , a subset (or projection) of the global logical dataset.

### 2.2 Operations

Clients issue **operations**  $op = (key, mutation, ctx)$ , where  $ctx$  carries metadata used for ordering and deduplication (Section 5). Operations are applied to local state first and recorded in a durable, ordered **change log**  $L_i$ . The synchronization protocol propagates entries of  $L_i$  to  $r_0$  and pulls remote changes back.

### 2.3 Channel

The channel between any  $r_i$  and  $r_0$  is **asynchronous and unreliable**: messages may be delayed unboundedly, dropped, duplicated, or reordered, and the channel may be unavailable (partitioned) for arbitrary durations. We do *not* assume a synchronized global clock; clocks are local, monotonic where the OS permits, and subject to skew and jumps.

### 2.4 Failure modes we design against

These are not edge cases. Each occurs routinely in a deployed fleet.

Failure	Description	Designed response
<b>Partition</b>	Device offline for seconds to weeks.	Local-first operation; queued change log; delta sync on reconnect.
<b>Process death</b>	OS kills the app, possibly mid-write.	Write-ahead durability; resumable work from persisted state.
<b>Power loss mid-write</b>	Battery dies or device powers off during a storage write.	Atomic/transactional writes; no partial- commit visibility.
<b>Message duplication</b>	Retried request arrives more than once.	Idempotency keys; server-side dedup window.
<b>Message reorder</b>	Operations arrive out of causal order.	Version vectors / causal metadata; reorder buffer.
<b>Concurrent update</b>	Two replicas mutate the same key while partitioned.	Conflict detection (not just resolution) via version vectors; typed merge.
<b>Server region outage</b>	Backend zone unavailable.	Client circuit breaker + cached fallback; backend multi-AZ.
<b>Clock skew/ jump</b>	Device wall clock wrong or non- monotonic.	Avoid wall-clock ordering for correctness; use logical clocks.
<b>Thundering herd</b>	Fleet reconnects/retries simultaneously after an outage.	Full-jitter backoff; server rate limiting.

A correct architecture has a defined, tested behavior for every row. The most common production incidents come from rows the team never wrote down.

### 3. CONSISTENCY MODELS AT THE MOBILE TIER

"Eventually consistent" is too coarse to design against. We need to commit to specific guarantees and know the conditions under which each holds.

- **Linearizability** (single-key strong consistency): every read returns the most recent committed write. Unachievable during partition without sacrificing availability the CAP trade-off [4, 5]. On mobile, requiring it for ordinary operations is a design error, because partition is the normal state.
- **Causal consistency**: if operation  $A$  happens-before  $B$  (Lamport's  $\rightarrow$ ) [1], every replica that observes  $B$  has already observed  $A$ . This is achievable while available under partition and is strong enough for most application semantics. It is the target for the protocol in Section 5.
- **Read-your-writes / monotonic-reads / monotonic-writes**: session guarantees [3].

#### 4.1 Durable, transactional local store (addresses process death, power loss)

The local store must give atomicity and durability for a write before the UI acknowledges it. The standard implementation is a write-ahead log (WAL): the mutation is appended to an append-only log and fsync'd, then applied to the main store; on restart, the log is replayed to recover any write that had not yet been folded in. The critical invariant is the **commit ordering**:

1. Append mutation to WAL, fsync. // now durable
2. Acknowledge to the user / UI. // safe to claim success
3. Enqueue for synchronization. // background

In a local-first design these come *for free* on the originating device, because the device reads from its own state, which always reflects its own writes. This is an underrated reliability benefit of local-first.

- **Eventual consistency** [10]: all replicas converge given no new updates and eventual connectivity. This is the floor we guarantee globally; causal consistency is the stronger guarantee we layer on top.

Design rule: choose the guarantee **per data type**, not once for the whole app. A user's draft text wants read-your-writes and last-writer-wins on conflict. A shared counter wants a CRDT. A financial ledger entry may need to be funneled through  $r_0$  for a strongly serialized decision and must not be merged on-device at all.

### 4. RELIABILITY MECHANISMS

Each subsection states the mechanism and then the failure it addresses from the table in §2.4.

Acknowledging before step 1 is the single most common source of "the app said it saved but my data is gone." If the process dies between 1 and 2, recovery replays the WAL and the write survives the user simply sees it on next launch.

#### 4.2 Resum'able work across the OS lifecycle (addresses process death)

Mobile operating systems terminate processes and revoke background execution on their own schedule; in-memory progress is not a safe place to keep anything you cannot lose. Any multi-step task (a sync cycle, a large upload, a migration) must be modeled as a

state machine whose current state is persisted, so that after a kill it resumes from the last durable checkpoint rather than restarting or aborting. Boot-completion and process-restart paths should carry first-class test coverage; in a fleet they execute constantly even though they are rare on a developer's desk.

#### 4.3 Idempotent operations and effectively-once delivery (addresses duplication)

Exactly-once *delivery* over an unreliable channel is impossible. The achievable and correct goal is **at-least-once delivery + idempotent application = effectively-once effect**. Each operation carries a client-generated stable idempotency key (a UUID minted at operation creation, *not* at send time, so retries reuse it):

```
op.ctx.idempotency_key = uuid_v4() // minted once, at creation
```

The server keeps a dedup table keyed by **idempotency\_key** over a retention window longer than the maximum plausible retry horizon. On receiving an op whose key is

already recorded, it returns the prior result without re-applying the mutation. This is what makes aggressive client retries safe.

#### 4.4 Circuit breaking and graceful degradation (addresses server outage, latency)

A client that hangs waiting on a degraded dependency turns a backend problem into a frozen UI. A circuit breaker [9] is a small state machine in front of each remote dependency:

```
states: CLOSED (normal) -> OPEN (failing fast) -> HALF_OPEN (probing) -> CLOSED
```

```
CLOSED:    pass calls through. On rolling failure rate >= F_threshold  
           within window W, trip to OPEN.
```

```
OPEN:      fail fast (return cached/fallback) for cooldown T_open.  
           After T_open, move to HALF_OPEN.
```

```
HALF_OPEN: allow a limited number of probe calls. If they succeed,  
           reset to CLOSED. If any fail, return to OPEN.
```

Combined with per-call timeouts and cached fallbacks, this bounds the user-visible latency of any single failing dependency and prevents a slow service from consuming all client request slots.

#### 4.5 Observability that survives a crash (addresses all of the above, operationally)

Because the code runs on devices the developer never holds, the system must be able to reconstruct what happened remotely. Minimum: structured crash reports with stack traces; a ring buffer of breadcrumb events persisted to disk and uploaded on next launch

(so the events leading to a crash are not lost with the process); and metrics for the reliability invariants themselves (sync success rate, time-to-convergence, queue depth). The metric that matters is not "did it crash" but "can I reconstruct the sequence that led here from data the device sent me."

#### 5. THE SYNCHRONIZATION PROTOCOL

This section specifies the protocol's ordering, idempotency, conflict detection, and convergence semantics as explicit algorithms rather than prose.

##### 5.1 Causal metadata: version vectors

To detect concurrent updates (rather than silently overwriting), each replicated key carries a **version vector**  $V$  [2]: a map from replica id to a per-replica monotonic counter. Replica  $r_i$  incrementing its own entry on each local mutation lets any two versions be compared:

```
compare(Va, Vb):  
  if Va[k] == Vb[k] for all k           -> EQUAL  
  else if Va[k] <= Vb[k] for all k      -> Vb DOMINATES (Vb is newer)  
  else if Vb[k] <= Va[k] for all k      -> Va DOMINATES  
  else                                   -> CONCURRENT (true conflict)
```

The fourth case — neither dominates — is the only real conflict, and version vectors let us detect it precisely instead of guessing from timestamps. This is why we do **not** order

writes by wall-clock time: clock skew and non-monotonic jumps would silently lose data by mislabeling concurrent updates as ordered ones.

## 5.2 Delta synchronization with a cursor

Full-state sync does not scale (payload grows with  $D$ , not with the change rate). Each client tracks a **cursor** an opaque, server-issued position in the change stream and pulls only changes after it. Deletions are represented as **tombstones** so that a delete propagates as a change rather than as an absence (an absence is indistinguishable from "not yet synced").

```
SYNC CYCLE (client ri):
```

```
PUSH:
```

```
batch <- pending entries from local change log Li
for each op in batch: ensure op.ctx.idempotency_key is set
send (cursor_i, batch) to r0
on success(ack, new_changes, new_cursor):
    mark batch entries as confirmed in Li
    goto PULL-APPLY with new_changes, new_cursor
```

```
PULL-APPLY (apply remote changes):
```

```
for each remote op in new_changes (in causal order):
    local <- Si[op.key]
    switch compare(op.V, local.V):
        op.V DOMINATES      -> apply op (remote is newer)
        local.V DOMINATES  -> ignore (we already have newer)
        EQUAL               -> ignore (duplicate)
        CONCURRENT         -> resolve(op, local) // Section 5.3
    merge version vectors: Si[op.key].V <- pointwise-max(op.V, local.V)
cursor_i <- new_cursor
persist cursor_i and Si transactionally
```

The two transactional persistence points (confirming pushed entries, and storing the new cursor + applied changes) are what make the cycle resumable: a process death mid-cycle leaves the system in a consistent prior

state, and the next cycle re-derives the rest. Because pushes are idempotent (§4.3), re-sending an unconfirmed batch after a crash is safe.

### 5.3 Conflict resolution by data type

When **compare** returns CONCURRENT, the resolution strategy is chosen by the type of the data, declared in the schema:

- **LWW-Register** (e.g., a profile field): keep the write with the higher logical version, breaking remaining ties deterministically (e.g., by replica id) so all replicas pick the *same* winner. Note: the loser is still surfaced to observability, because silent loss of a user edit is a reliability defect even when convergence is technically achieved.
- **CRDTs** [6] for data that must merge without losing intent:
- *G-Counter* / *PN-Counter* counters that merge by per-replica max (and a negative partition for decrements). Merge:  $c[k] = \max(ca[k], cb[k])$  per replica.
- *OR-Set* (*observed-remove set*) add/remove set where each add carries a unique tag; remove only affects tags it has observed, so a concurrent add survives a concurrent remove. This avoids the "remove wins forever" pathology of naive sets.
- *RGA* / *sequence CRDTs* for collaboratively edited ordered text. CRDTs guarantee that **merge** is commutative, associative, and idempotent, which is exactly what makes convergence independent of message order and duplication the same properties our channel violates.
- **Server-arbitrated** for data that must not

be merged on-device at all (e.g., a balance decrement, an inventory claim): the client submits an *intent*, and  $r_o$  serializes the decision under strong consistency and returns accept/reject. The client treats its local optimistic state as provisional until confirmed.

### 5.4 Convergence

**Proposition.** Under the model of §2, assume (a) every operation is recorded in a durable per-replica log; (b) each replicated key carries a version vector and updates are applied per the PULL-APPLY rule of §5.2; (c) for each data type the resolution function is either a CRDT merge commutative, associative, and idempotent, satisfying the Strong Eventual Consistency conditions of Shapiro et al. [6] or a deterministic total order (LWW with a fixed tie-break, §5.3); and (d) connectivity is eventually restored for unbounded-but-finite periods so that every logged operation is eventually delivered to every replica at least once. Then all replicas converge to identical state for each key.

*Proof sketch.* By (d) every operation reaches every replica at least once. Duplicates are absorbed: version-vector comparison classifies a re-delivered op as EQUAL or dominated and ignores it, and CRDT merges are idempotent, so at-least-once delivery has the effect of exactly-once application (§4.3). Reordering is absorbed: the outcome of PULL-APPLY depends only on the version-vector partial order, not on arrival order,

and the merge functions are commutative

and associative, so any delivery order yields the same state this is exactly the Strong Eventual Consistency result of [6] for the CRDT-typed keys. For LWW-typed keys, the deterministic total order means every replica, once it has seen the same set of writes, selects the same winner; the version vector guarantees each replica eventually sees that set. Server-arbitrated keys (§5.3) are serialized at  $r_0$  and therefore trivially agree once the decision propagates. Hence every key converges. ■

The argument inherits the formal guarantee for the CRDT case directly from [6]; the contribution here is the composition rule that a schema mixing CRDT, LWW, and server-arbitrated keys still converges, because each key's type independently satisfies one of the two convergence conditions. A fully mechanized proof of the mixed scheme is left as future work (§10).

## 6. SCALABILITY AND THE QUANTITATIVE BEHAVIOR OF A FLEET

### 6.3 Retry and backoff: avoiding the self-inflicted DDoS

A fleet that retries on a fixed schedule will synchronize its retries and hammer the backend in waves most violently right after an outage ends, when  $N$  clients all reconnect at once (thundering herd). The mitigation is **exponential backoff with full jitter** [8]:

```
base = initial_delay
cap = max_delay
attempt n -> sleep = random_uniform(0, min(cap, base * 2^n))
```

### 6.1 Stateless, horizontally scalable backend

The server replica  $r_0$  is logically singular but physically a stateless, horizontally scaled service: any instance can handle any client's sync cycle because session and dedup state live in shared stores (a distributed cache + the change-stream store), not in instance memory. This is the standard shape of a highly available, partition-tolerant backing store [11]. It lets capacity track  $N$  by adding instances behind a load balancer, and lets the dedup table (§4.3) be a shared, TTL-bounded structure.

### 6.2 Load is sublinear in data because sync is delta-based

Per the cursor protocol (§5.2), the bytes a client exchanges per cycle are proportional to the number of changes since its cursor, not to  $D$ . So aggregate backend egress scales with the fleet's *change rate*, not its *data size* the property that lets per-user data grow without backend cost growing with it.

Full jitter (sampling uniformly from  $[0, \text{window}]$  rather than using the window itself) is what de-correlates the fleet: instead of  $N$  clients retrying at  $t = \text{base} \cdot 2^n$  together, their retries spread uniformly across the interval, flattening the peak request rate from  $\sim N$  simultaneous to  $\sim N / \text{window}$  per unit time. The server complements this with rate limiting and, on overload, returning a **Retry-After** that client's honor as a floor on their jitter window.

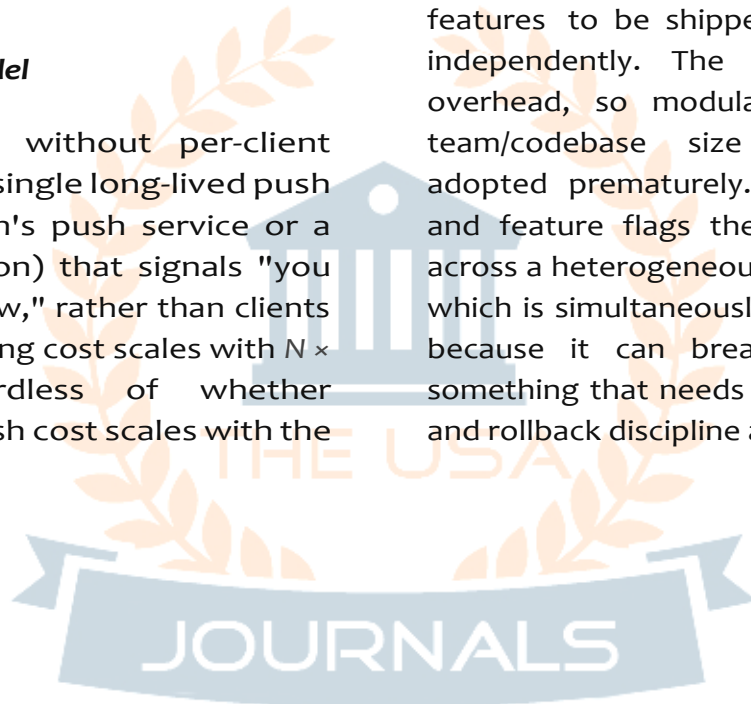
#### 6.4 Connection model

For push freshness without per-client polling cost, prefer a single long-lived push channel (the platform's push service or a multiplexed connection) that signals "you have changes, sync now," rather than clients polling on a timer. Polling cost scales with  $N \times \text{poll\_frequency}$  regardless of whether anything changed; push cost scales with the

change rate. This mirrors the delta-vs-dump argument at the connection layer.

#### 6.5 Client codebase scalability

Scalability also means the codebase and team can grow. Feature modules with enforced boundaries (a module may depend on published interfaces of another, not its internals) let teams work in parallel, keep incremental build times bounded, and allow features to be shipped, flagged, or disabled independently. The cost is real up-front overhead, so modularization should track team/codebase size rather than being adopted prematurely. Remote configuration and feature flags then let behavior change across a heterogeneous fleet without a release which is simultaneously a scalability lever and, because it can break the fleet remotely, something that needs the same staged-rollout and rollback discipline as code.





local store with keys held in the platform's hardware-backed keystore (Secure Enclave / StrongBox / TEE), not in app storage; bind key availability to device unlock where the data class warrants it.

- **Sync channel.** The delta channel is a target for interception, replay, and tampering. Mitigation: mutual TLS or token-bound transport; per-message authentication; and replay protection that reuses the idempotency key + a server-side nonce/window so a replayed op is recognized and discarded (the same machinery that gives effectively-once also resists replay).
- **Malicious or compromised client.** A client can submit forged operations. Server-side authorization must be enforced on  $r_0$  per operation; the optimistic local apply is a UX convenience, never an authorization decision. Server-arbitrated data types (§5.3) are exactly where this matter.
- **Lost/stolen device.** Mitigation: remote wipe via the config/flag channel, short-lived sync tokens, and encryption keys that become unusable after remote revocation.
- **Clock manipulation.** Because correctness uses logical clocks (version vectors), an attacker setting the device clock cannot reorder or win conflicts a side benefit of not ordering by wall time.

## 9. EVALUATION FRAMEWORK AND WORKED EXAMPLE

### 9.1 Metrics (each tied to a §2 property)

**Reliability** - *Crash-free session rate and crash-free user rate. - Durable-write loss incidents per release target zero; measured by reconciling acknowledged writes against converged state. - Sync success rate = successful cycles / attempted cycles. - Time-to-convergence (TTC) after reconnection = wall-clock from connectivity restored to all pending deltas applied and acknowledged. Report distribution (p50/p95/p99), not just mean. - Conflict rate and silent-loss rate concurrent updates per 1k ops, and how many resulted in an unrecoverable user-edit loss (should be zero by design).*

**Scalability** - *Backend cost per monthly active client as  $N$  grows should be flat/ declining. - Sync bytes per operation and its ratio to  $D$  should track change rate, not data size. - Peak request rate after simulated outage with vs. without full-jitter backoff - Incremental build time and merge-conflict rate as codebase/team grow.*

### 9.2 Fault-injection methodology

Reliability claims are only credible if tested under the failure model, not under normal conditions. The harness should be able to:

- kill the process at randomized points, including mid-WAL-write and mid-sync-cycle, and assert no durable-write loss and successful resume;
- partition the channel for randomized durations and assert offline functionality + post-reconnect convergence;

- duplicate, drop, delay, and reorder sync messages and assert idempotent/causal correctness;
- inject server 5xx/latency and assert circuit-breaker behavior and bounded UI latency;
- skew/jump the device clock and assert conflict outcomes are unchanged;
- simulate  $N$ -client simultaneous reconnect and measure peak server load.

Failures should be reproducible: driving the harness from a deterministic seed, and adopting a property-based or simulation approach (in the spirit of deterministic simulation testing), allows the same failure scenario to be replayed and bisected rather than observed once and lost.

### 9.3 Worked example: a consumer cross-device application

To make the framework concrete, consider a representative deployment: a consumer productivity application (notes, tasks, and settings) used by a large installed base, where each user runs it across several of their own gadgets a phone, a tablet, and a wearable and edits the same data from any of them. This stresses every part of the model. Partitions are long and routine (airplane mode, subway and elevator dead zones, a phone left off overnight, a watch out of Bluetooth range of its phone). Process and power cycles are frequent (the OS evicts backgrounded apps under memory pressure, and small gadgets sleep aggressively to save battery).

Concurrent updates are the common case rather than the exception, because a single user genuinely edits the same note on their phone and tablet while both are briefly offline. And the installed base is large enough that uncoordinated retries are a real self-DDoS risk after any backend hiccup.

The architecture maps cleanly. A note or task edit is local-first: it commits to the WAL and is acknowledged instantly, so the UI never blocks on the network and reads-your-writes holds on whichever gadget made the edit. Text fields use LWW-Registers, ordered lists and tag sets use sequence/OR-Set CRDTs so a concurrent add on the phone is not lost to a concurrent remove on the tablet, and the version vector detects the genuinely concurrent edits between a user's own devices rather than letting a stale write from a just-reconnected wearable clobber newer state. Account-level settings that must be unambiguous (e.g., subscription entitlement) are server-arbitrated. Long offline periods are absorbed by the durable change log and replayed on reconnect; OS-initiated process death mid-sync is handled by the resumable, WAL-backed cycle. Full-jitter backoff matters here specifically because millions of clients can re-attempt sync in the same window when a backend zone recovers, and uncoordinated retries would turn recovery into a second outage.

This example is presented as a design analysis rather than a measured study: it shows how each failure mode in Section 2 is discharged by a specific mechanism, and which metrics of Section 9.1 would be instrumented in a

deployment. Reporting those metrics from a production system is the natural next step, and is the strongest form of validation for the architecture.

## 10. CHALLENGES AND FUTURE DIRECTIONS

- **Security vs. local-first.** Reconciling an authoritative on-device replica with strong confidentiality, especially for regulated data, remains under-specified in practice.
- **On-device models as state.** As inference moves on-device, models become versioned, updatable state that the data/sync layers were not designed to manage (size, fallback, staged rollout of weights).
- **Energy as a first-class budget.** Retries, background sync, encryption, and observability all cost energy; architectures that budget joules the way they budget memory are not yet standard.
- **Long-horizon compatibility.** A large installed base typically has many app versions live at once, and gadgets can sit unused for long stretches and then reconnect on old software; the wire protocol and schema must stay forward/ backward-compatible across a wide version span, and schema/protocol evolution over that horizon is under-theorized.
- **Formalizing the convergence guarantee** of §5.4 with mechanized proof, rather than sketch, for the specific mixed LWW/CRDT/server-arbitrated scheme used

here.

## 11. CONCLUSION

Treating the mobile client as a first-class replica in an explicitly modeled distributed system rather than as a view onto a backend turns the field's recurring failures into problems with known solutions. From a concrete system and failure model we derived per-type consistency guarantees, a delta synchronization protocol with idempotent, causally-ordered, conflict-detecting semantics, and a quantitative account of why full-jitter backoff and delta sync keep a large fleet from overwhelming its own backend. The same local-first foundation delivers reliability under partition and scalability of backend cost, because the mechanisms coincide. The evaluation framework and fault-injection methodology are included to keep the claims falsifiable; the open problems mark where the next work lies.

## REFERENCES

1. L. Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Communications of the ACM*, vol. 21, no. 7, pp. 558–565, Jul. 1978. doi: 10.1145/359545.359563
2. D. S. Parker, G. J. Popek, G. Rudisin, A. Stoughton, B. J. Walker, E. Walton, J. M. Chow, D. Edwards, S. Kiser, and C. Kline, "Detection of Mutual Inconsistency in Distributed Systems," *IEEE Transactions on Software Engineering*, vol. SE-9, no. 3, pp. 240–247, May 1983. doi:10.1109/TSE.1983.236733
3. D. B. Terry, A. J. Demers, K. Petersen, M.

- J. Spreitzer, M. M. Theimer, and B. B. Welch, "Session Guarantees for Weakly Consistent Replicated Data," in *Proc. 3rd Int. Conf. on Parallel and Distributed Information Systems (PDIS)*, Austin, TX, USA, Sep. 1994, pp. 140–149. doi:10.1109/PDIS.1994.331722
4. S. Gilbert and N. Lynch, "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services," *ACM SIGACT News*, vol. 33, no. 2, pp. 51–59, Jun. 2002. doi:10.1145/564585.564601
  5. E. Brewer, "CAP Twelve Years Later: How the 'Rules' Have Changed," *Computer*, vol. 45, no. 2, pp. 23–29, Feb. 2012. doi:10.1109/MC.2012.37
  6. M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski, "Conflict-Free Replicated Data Types," in *Stabilization, Safety, and Security of Distributed Systems (SSS 2011)*, LNCS, vol. 6976. Berlin, Heidelberg: Springer, 2011, pp. 386–400. doi:10.1007/978-3-642-24550-3\_29
  7. M. Kleppmann, A. Wiggins, P. van Hardenberg, and M. McGranaghan, "Local-First Software: You Own Your Data, in spite of the Cloud," in *Proc. 2019 ACM SIGPLAN Int. Symp. on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! '19)*, Athens, Greece, 2019, pp. 297–307. doi:10.1145/3359591.3359737
  8. M. Brooker, "Exponential Backoff And Jitter," AWS Architecture Blog, Mar. 2015. [Online]. Available: <https://aws.amazon.com/blogs/architecture/exponential-backoff-and-jitter/>
  9. M. Nygard, *Release It!: Design and Deploy Production-Ready Software*, 2nd ed. Raleigh, NC, USA: Pragmatic Bookshelf, 2018.
  10. W. Vogels, "Eventually Consistent," *Communications of the ACM*, vol. 52, no. 1, pp. 40–44, Jan. 2009. doi:10.1145/1435417.1435432
  11. G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: Amazon's Highly Available Key-value Store," in *Proc. 21st ACM SIGOPS Symp. on Operating Systems Principles (SOSP '07)*, Stevenson, WA, USA, 2007, pp. 205–220. doi:10.1145/1294261.1294281
  12. R. C. Martin, *Clean Architecture: A Craftsman's Guide to Software Structure and Design*. Boston, MA, USA: Prentice Hall, 2017.