



**HAL**  
open science

# Write Fast, Read in the Past: Causal Consistency for Client-side Applications

Marek Zawirski, Nuno Pregoça, Sérgio Duarte, Annette Bieniusa, Valter Balegas, Marc Shapiro

## ► To cite this version:

Marek Zawirski, Nuno Pregoça, Sérgio Duarte, Annette Bieniusa, Valter Balegas, et al.. Write Fast, Read in the Past: Causal Consistency for Client-side Applications. [Research Report] RR-8729, Inria – Centre Paris-Rocquencourt; INRIA. 2015. hal-01158370v1

**HAL Id: hal-01158370**

**<https://inria.hal.science/hal-01158370v1>**

Submitted on 31 May 2015 (v1), last revised 15 Dec 2015 (v2)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# Write Fast, Read in the Past: Causal Consistency for Client-side Applications

**Marek Zawirski** Inria & Sorbonne Universités, UPMC Univ Paris 06, LIP6

**Nuno Preguiça** NOVA-LINCS, U. Nova de Lisboa

**Sérgio Duarte** NOVA-LINCS, U. Nova de Lisboa

**Annette Bieniusa** U. Kaiserslautern

**Valter Balegas** NOVA-LINCS, U. Nova de Lisboa

**Marc Shapiro** Inria & Sorbonne Universités, UPMC Univ Paris 06, LIP6

**RESEARCH  
REPORT**

**N° 8729**

May 2015

Project-Team Regal





## Write Fast, Read in the Past: Causal Consistency for Client-side Applications

Marek Zawirski Inria & Sorbonne Universités, UPMC Univ Paris 06, LIP6

Nuno Preguiça NOVA-LINCS, U. Nova de Lisboa

Sérgio Duarte NOVA-LINCS, U. Nova de Lisboa

Annette Bieniusa U. Kaiserslautern

Valter Balegas NOVA-LINCS, U. Nova de Lisboa

Marc Shapiro Inria & Sorbonne Universités, UPMC Univ Paris 06, LIP6

Project-Team Regal

Research Report n° 8729 — May 2015 — 29 pages

**Abstract:** Client-side apps (e.g., mobile or in-browser) need cloud data to be available in a cache, for both reads and updates. The cache should use resources sparingly, be consistent and fault-tolerant. The system needs to scale to high numbers of unreliable and resource-poor clients, and large database. The SwiftCloud distributed object database is the first to provide fast reads and writes via a causally-consistent client-side local cache. It is thrifty in resources and scales well, thanks to consistent versioning provided by the cloud, using small and bounded metadata. It remains available during faults, switching to a different data centre when the current one is not responsive, while maintaining its consistency guarantees. This paper presents the SwiftCloud algorithms, design, and experimental evaluation. It shows that client-side apps enjoy the high performance and availability, under the same guarantees as a remote cloud data store, at a small cost.

**Key-words:** geo-replication, causal consistency, eventual consistency, fault tolerance, client-side applications

**RESEARCH CENTRE  
PARIS – ROCQUENCOURT**

Domaine de Voluceau, - Rocquencourt  
B.P. 105 - 78153 Le Chesnay Cedex

## Écriture rapide, lecture dans le passé : cohérence causale pour les applications localisées côté client

**Résumé :** Les applications localisées côté client, par exemple les applications mobiles ou les applications dans le navigateur, ont besoin que les données du nuage soient disponibles dans un cache local, aussi bien pour l'écriture que pour la lecture. Ce cache doit utiliser les ressources avec économie, être cohérent, et tolérer les fautes. Le système doit passer à l'échelle d'un nombre élevé de clients non fiables et disposant de peu de ressources, et d'une base de données de grande capacité. La base de données répartie SwiftCloud est la première qui permet les lectures et les écritures rapides et cohérentes par l'intermédiaire d'un cache situé côté client. Elle utilise les ressources avec parcimonie, et passe bien à l'échelle, grâce au fait que le nuage sert des versions cohérentes tout en ne nécessitant que des métadonnées de taille petite et bornée. Les données restent disponibles pendant les fautes de réseau, car, quand le centre de calcul courant ne répond pas, le cache peut passer à un nouveau centre de calcul sans violer ses garanties de cohérence.

Le présent papier présente les algorithmes, la conception et l'évaluation expérimentale de SwiftCloud. Il montre que le système fournit aux applications côté client, à la fois des performances et une disponibilité élevées, et à la fois les mêmes garanties qu'un stockage distant dans le nuage, à un faible coût.

**Mots-clés :** géo-réplication, cohérence causale, cohérence finale, tolérance aux pannes, applications localisées côté client

## 1 Introduction

Client-side applications, such as in-browser and mobile apps, are poorly supported by the current technology for sharing mutable data over the wide area. App developers end up resorting to ad-hoc application-level caching and buffering implementations, in order to avoid slow, costly and sometimes unavailable round-trips to a data centre, but they cannot solve system issues such as fault tolerance or session guarantees [37]. Recent application frameworks such as Google Drive Realtime API [14], TouchDevelop [12] or Mobius [15] support client-side access at a small scale, but do not provide system-wide consistency and/or fault tolerance guarantees. Algorithms for geo-replication [4, 6, 19, 27, 28] or for managing database replicas on clients [9, 30] ensure some of the right properties, but were not designed to support high numbers of client replicas.

Our thesis is that it is up to the system to provide client-side applications with correct and scalable database access. The system should address the (somewhat conflicting) requirements of consistency, availability, and convergence [29], and uphold them at least as well as a server-side system. Concurrent updates (which are unavoidable if updates are to be always available) should neither be lost nor cause the database to diverge permanently. Under these requirements, the strongest possible consistency model is *causal consistency with convergent objects* (CRDTs) [27, 29, 35].

Supporting thousands or millions of client-side replicas challenges classical assumptions, as we explain in more detail later in this article: (i) To track causality precisely, per client, creates unacceptably fat metadata; but the more compact server-side metadata management has fault-tolerance issues. (ii) Full replication at high numbers of resource-poor devices would be unacceptable [9]; but partial replication of data and metadata could cause anomalous message delivery or unavailability. (iii) Unlike many previous approaches [4, 19, 27, 28], fault tolerance and consistency cannot be solved by assuming that the application is located inside the data centre (DC), or has a sticky session to a single DC [7, 37].

This work addresses these challenges. We present the algorithms, design, and evaluation of SwiftCloud, the first distributed object store designed for a high number of replicas. It efficiently ensures consistent, available, and convergent access to client nodes, tolerating failures. To enable both small metadata and fault tolerance, SwiftCloud uses a flexible client-server topology, and decouples reads from writes. The client *writes fast* into the local cache, and *reads in the past* (also fast) data that is consistent, but occasionally stale. Our main contribution is a combination of two techniques:

**Cloud-backed support for partial replicas** (Section 3) A DC serves a consistent view of the database to the client, which the client merges with its own updates. In some failure situations, a client may connect to a DC that happens to be inconsistent with its previous DC. Because the client does not have a full replica, it cannot fix the issue on its own. We leverage “reading in the past” to avoid this situation in the common case, and provide control over the inherent trade-off between staleness and unavailability. More precisely, a client observes a remote update only if it is stored in some number  $K \geq 1$  of DCs [30]; the higher the value of  $K$ , the more likely that a  $K$ -stable version is in both DCs, but the higher the staleness.

**Protocols with decoupled, bounded metadata** (Section 4) Thanks to funnelling communication through DCs and to “reading in the past,” SwiftCloud leverages decoupled metadata [25] separating *tracking causality*, which uses small vectors assigned in the background by DCs, from *unique identification*, based on client-assigned scalar timestamps. Consequently, the metadata is small and bounded in size. Furthermore, a DC can prune its log independently of clients, replacing it with a summary of delivered updates.

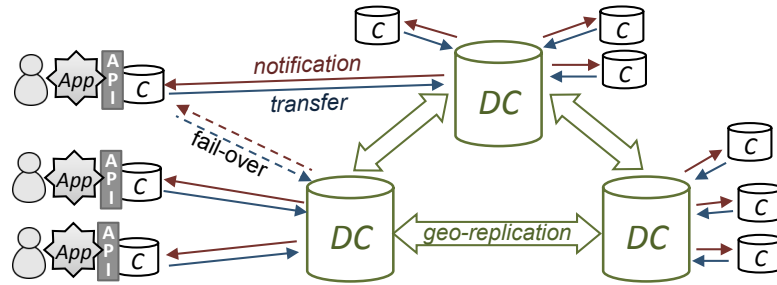


Figure 1: System components (Application processes, Clients, Data Centres), and their interfaces.

We implement SwiftCloud and demonstrate experimentally that our design reaches its objective, at a modest staleness cost. We evaluate SwiftCloud in Amazon EC2, against a port of WaltSocial [36] and against YCSB [16]. When data is cached, response time is two orders of magnitude lower than for server-based protocols with similar availability guarantees. With three DCs (servers), the system can accommodate thousands of client replicas. Metadata size does not depend on the number of clients, the number of failures, or the size of the database, and increases only slightly with the number of DCs: on average, 15 bytes of metadata overhead per update, compared to kilobytes for previous algorithms with similar safety guarantees. Throughput is comparable to server-side replication, and improved for high locality workloads. When a DC fails, its clients switch to a new DC in under 1000 ms, and remain consistent. Under normal conditions, 2-stability causes fewer than 1% stale reads.

The paper is organised as follows. In Section 2, we state the problem by motivating the system model and the requirements of client-side database access. In Section 3, we present the design principles of SwiftCloud, and its concrete protocols in Section 4. In Section 5 we present our experimental results, followed by discussion of related work in Section 6.

## 2 Problem overview

We consider support for a variety of client-side applications, sharing a database of *objects* that clients can read and update. We aim to scale to thousands of clients, spanning the whole internet, and to a database of arbitrary size.

Figure 1 illustrates our system model. A cloud infrastructure connects a small set (say, tens) of geo-replicated data centres, and a large set (thousands) of clients. A DC has abundant computational, storage and network resources. Similarly to Sovran et al. [36], we abstract a DC as a powerful sequential process that hosts a *full replica* of the database.<sup>1</sup> DCs communicate in a peer-to-peer way. A DC may fail (e.g., due to disaster, power outage, WAN partition or misconfiguration [5, 24]) and recover with its persistent memory intact.

Clients do not communicate directly, but only via DCs. Normally, a client connects to a single DC; in case of failure or roaming, to zero or more. A client may fail and recover (e.g., disconnection while travelling) or fail permanently (e.g., destroyed phone), both without prior

<sup>1</sup>We refer to prior work for the somewhat orthogonal issues of parallelism and fault-tolerance within a DC [4, 19, 27, 28].

warning. We consider only non-byzantine failures.<sup>2</sup>

Client-side apps need to respond quickly and at all times, i.e., they require high *availability* and *responsiveness*. This can be achieved by replicating data locally, and by synchronising updates in the background. However, a client has limited resources; therefore, it hosts a *cache* that contains only the small subset of the database of current interest to the local app. It should not have to receive messages relative to objects that it does not currently replicate [33]. Finally, control messages and piggy-backed metadata should have small and bounded size.

Since a client replica is only *partial*, there cannot be a guarantee of complete availability. The best we can expect is *partial availability*, whereby an operation returns without remote communication if the requested data is cached; and after retrieving the data from a remote node (DC) if not. If the data is not there and the network is down, the operation may be unavailable, i.e., it either blocks or returns an error.

## 2.1 Consistency with convergence

Application programmers and users wish to observe a consistent view of the global database. However, with availability as a requirement, consistency options are limited [21, 29].

**Causal consistency** The strongest available and convergent model is causal consistency [2, 29].

**Definition 1 (Causal order, causal consistency)** *Let an execution  $E$  be a set of sequences (one per application process) of operation invocations with their return values. Operations  $a$  and  $b$  are potentially causally-related in  $E$ , noted  $a \rightarrow b$  and called causal order, if:*

1. *An application process invoked  $b$  after it invoked  $a$ ; or*
2. *Read  $b$  observed update  $a$  on the same object; or*
3. *There exists an operation  $c \in E$  such that  $a \rightarrow c \rightarrow b$ .*

*An execution  $E$  is causally consistent if every read operation in  $E$  observes all the updates that causally precede the read, applied in some linear extension of causal order.*

Informally, under causal consistency, every process observes a *monotonically non-decreasing set of updates that includes its own updates, in an order that respects the causality between operations*.<sup>3</sup> The following well-known example illustrates [27]. In a social network, Bob sets permissions to disallow his boss Alice from viewing his photos. Some time later, Bob posts a questionable photo of himself. Without causal consistency, Alice may view the bad photo, delivered before the new permissions. Under causal consistency, the change of permission is guaranteed to be delivered before the post, and Alice cannot view the photo.

More generally, if an application process reads  $x$ , and later reads  $y$ , and if the state of  $x$  causally-depends on some update  $u$  to  $y$ , then the state of  $y$  that it reads will include update  $u$ . When the application requests  $y$ , we say there is a *causal gap* if the local replica has not yet received  $u$ . A consistent system must detect such a gap, and wait until  $u$  is delivered before returning  $y$ , or avoid it in the first place. If not, inconsistent reads expose both programmers and users to anomalies caused by gaps [27, 28].

We extend causal consistency with multi-operation *causal transactions*. All the reads of such a transaction come from a same database *snapshot*, and either all its updates are visible as a

<sup>2</sup>Possible scalable approaches for Byzantine clients include admission control, and the Fork Join Causal Consistency protocol [30] to protect metadata, at the DC perimeter.

<sup>3</sup>This subsumes the well-known session guarantees [13].



group, or none is [27, 28]. The snapshot is causally consistent, and the transaction’s updates are considered to causally depend on all of the transaction’s reads.

**Convergence** Applications require *convergence*, which consists of liveness and safety properties: (i) *At-least-once delivery*: an update that is delivered (i.e., is visible by the app) at some node, is delivered to all interested nodes<sup>4</sup> after a finite number of message exchanges; (ii) *Confluence*: nodes that delivered the same set of updates read the same value.

Causal consistency does not guarantee confluence, as two replicas might receive the same updates in different orders. For confluence, we rely on CRDTs, high-level data types with rich confluent semantics [13, 35]. An update on a high-level object is not just an assignment, but is a method associated with the object’s type. For instance, a Set object supports `add(element)` and `remove(element)`; a Counter supports `increment()` and `decrement()`.

CRDTs include primitive last-writer-wins register (LWW) and multi-value register (MVR) [18, 23], but also higher level types such as Sets, Lists, Maps, Graphs, Counters, etc. [1, 34–36]. Registers are simple to implement [19, 27, 28], but cumbersome to use. For instance, the implementation of LWW needs to store only the “last” update and has idempotent updates, but loses some concurrent assignments. Higher level types, such as Counter or Set, do not lose updates and are easier to use, but their implementation is more demanding. The implementation of high-level objects is eased by adequate support from the system. For instance, an object’s value may be defined not just by the last update, but also depend on earlier updates; causal consistency is essential, ensuring that they are not lost nor delivered out of order. Safety also demands *at-most-once delivery*, as high-level updates are often not idempotent (consider for instance `increment()`).

Although each of these requirements may seem familiar or simple in isolation, the combination with scalability to high numbers of nodes and database size is a new challenge.

## 2.2 Metadata design

*Metadata* serves to identify updates and to ensure correct delivery. Metadata is piggy-backed on update messages, increasing the cost of communication.

One common metadata design assigns each update a timestamp as soon as it is generated on some originating node. The causality data structures tend to grow “fat.” For instance, dependency lists [27] grow with the number of updates [19, 28], whereas version vectors [9, 30] grow with the number of clients. (Indeed, our experiments hereafter show that their size becomes unreasonable). We call this the *Client-Assigned, Safe but Fat* approach.

An alternative delegates timestamping to a small number of DC servers [4, 19, 28]. This enables the use of small vectors, at the cost of losing some parallelism. However, this is not fault tolerant if the client does not reside in a DC failure domain. For instance, it may violate *at-most-once delivery*. Consider a client transmitting update  $u$  to be timestamped by  $DC_1$ . If it does not receive an acknowledgement, it retries, say with  $DC_2$  (failover). This may result in  $u$  receiving two distinct timestamps, and being delivered twice. Duplicate delivery violates safety for many confluent types, or otherwise complicates their implementation [3, 13, 28]. We call this the *Server-Assigned, Lean but Unsafe* approach.

Clearly, neither “fat” nor “unsafe” is satisfactory.

<sup>4</sup> Interested nodes are those that replicate or cache the updated object.

### 2.3 Causal consistency with partial replication is hard

Since a partial replica receives only a subset of the updates, and hence of metadata, it could miss some causal dependencies [9]. Consider the following example: Alice posts a photo on her wall in a social network application (update **a**). Bob sees the photo and mentions it in a message to Charles (update **b**), who in turn mentions it to David (update **c**). When David looks at Alice’s wall, he expects to observe update **a** and view the photo. However, if David’s machine does not cache Charles’ inbox, it cannot observe the causal chain  $\mathbf{a} \rightarrow \mathbf{b} \rightarrow \mathbf{c}$  and might incorrectly deliver **c** without **a**. Metadata design should protect from such causal gaps, caused by transitive dependency over absent objects.

Failures complicate the picture even more. Suppose David sees Alice’s photo, and posts a comment to Alice’s wall (update **d**). Now a failure occurs, and David’s machine fails over to a new DC. Unfortunately, the new DC has not yet received Bob’s update **b**, on which comment **d** causally depends. Therefore, it cannot deliver the comment, i.e., fulfill convergence, without violating causal consistency. David cannot read new objects from the DC for the same reason.<sup>5</sup>

Finally, a DC logs an individual update for only a limited amount of time, but clients may be unavailable for unlimited periods. Suppose that David’s comment **d** is accepted by the DC, but David’s machine disconnects before receiving the acknowledgement. Much later, after **d** has been executed and purged away, David’s machine comes back, only to retry **d**. This could violate at-most-once delivery; some previous systems avoid this with fat version vectors [9, 30] or depend on client availability [25].

## 3 The SwiftCloud approach

We now describe an abstract design that addresses the above challenges, first in the failure-free case, and next, how we support DC failure. Our design applies the principles of consistency algorithms for full replication systems to build a cloud-based support for partial client replicas. Later, in Section 4, we present a concrete protocol implementing our design.

### 3.1 Causal consistency at full DC replicas

Ensuring causal consistency at fully-replicated DCs is a well-known problem [2, 19, 27, 28]. Our design is log-based, i.e., SwiftCloud stores updates in a log and transmits them incrementally; it includes optimisations, where the full log is occasionally replaced by the state of an object, called *checkpoint* [9, 31]. We discuss checkpoints only where relevant.

A *database version* is any subset of updates, noted  $U$ , ordered by causality. A version maps object identifiers to object state, by applying the relevant subsequence of the log; the value of an object is exposed via the read API.

We say that a version  $U$  has a *causal gap*, or is *inconsistent* if it is not causally-closed, i.e., if  $\exists u, u' : u \rightarrow u' \wedge u \notin U \wedge u' \in U$ . As we illustrate shortly, reading from an inconsistent version should be avoided, because, otherwise, subsequent accesses might violate causality. On the other hand, waiting for the gap to be filled would increase latency and decrease availability. To side-step this conundrum, we adopt the approach of “reading in the past” [2, 27]. Thus, a DC exposes a gapless but possibly delayed state, noted  $V$ .

To illustrate, consider the example of Figure 2(a). Objects  $x$  and  $y$  are of type `Set`.  $DC_1$  is in state  $U_1$  that includes version  $V_1 \subseteq U_1$ , and  $DC_2$  in a later state  $V_2$ . Versions  $V_1$  with

<sup>5</sup>Note that David can still perform updates, but they cannot be delivered, thus the system does not converge.




Figure 2: Example evolution of configurations for two DCs, and a client.  $x$  and  $y$  are Sets; box = update; arrow = causal dependence (an optional text indicates the source of dependency); dashed box = named database version/state.

value  $[x \mapsto \{1\}, y \mapsto \{1\}]$  and  $V_2$  with value  $[x \mapsto \{1, 3\}, y \mapsto \{1, 2\}]$  are both gapless. However, version  $U_1$ , with value  $[x \mapsto \{1, 3\}, y \mapsto \{1\}]$  has a gap, missing update  $y.add(2)$ . When a client requests to read  $x$  at  $DC_1$  in state  $U_1$ , the DC could return the most recent version,  $x : \{1, 3\}$ . However, if the application later requests  $y$ , to return a safe value of  $y$  requires to wait for the missing update from  $DC_2$ . By “reading in the past” instead, the same replica exposes the older but gapless version  $V_1$ , reading  $x : \{1\}$ . Then, the second read will be satisfied immediately with  $y : \{1\}$ . Once the missing update is received from  $DC_2$ ,  $DC_1$  may advance from version  $V_1$  to  $V_2$ .

A gapless algorithm maintains a *causally-consistent, monotonically non-decreasing progression* of replica states [2]. Given an update  $u$ , let us note  $u.deps$  its set of causal predecessors, called its *dependency set*. If a full replica, in some consistent state  $V$ , receives  $u$ , and its dependencies are satisfied, i.e.,  $u.deps \subseteq V$ , then it applies  $u$ . The new state is  $V' = V \oplus \{u\}$ , where we note by  $\oplus$  the *log merge operator* that unions the two logs, respecting causality (the operator filters out duplicates, as we discuss in Section 4.1). State  $V'$  is consistent, and monotonicity is respected, since  $V \subseteq V'$ .

If the dependencies are not met, the replica buffers  $u$  until the causal gap is filled.

### 3.2 Causal consistency at partial client replicas

As a client replica contains only part of the database and its metadata, this complicates consistency [9]. To relieve clients of this complexity, we move most of the burden of managing gapless versions to the DC-side full replica.

At any point in time, a client is interested in a subset of the objects in the database, called its *interest set*. Its initial state consists of the projection of some DC's state onto its interest set. This is a causally-consistent state, as shown in the previous section. Client state can change, either because of an update generated by the client itself, called an *internal update*, or because of one received from a DC, called *external*. An internal update obviously maintains causal consistency. If the same DC sends external updates with no gaps, then the client state remains causally consistent.

More formally, let  $i \in \mathcal{DC}$  denote DC identifiers and  $m \in \mathcal{C}$  client identifiers with  $\mathcal{DC} \cap \mathcal{C} = \emptyset$ .<sup>6</sup>

Consider some recent DC state, which we will call the *base version* of the client, noted  $V_{DC}$ . The interest set of client  $C$  is noted  $O \subseteq x, y, \dots$ . The client state, noted  $V_C$ , is restricted to these objects. It consists of two parts. One is the projection of base version  $V_{DC}$  onto its interest set, noted  $V_{DC}|_O$ . The other is the log of internal updates, noted  $U_C$ . The client state is their merge  $V_C = V_{DC}|_O \oplus U_C|_O$ . On cache miss, the client adds the missing object to its interest set, and fetches the object from base version  $V_{DC}$ , thereby extending the projection.

Base version  $V_{DC}$  is a monotonically non-decreasing causal version (it might be slightly behind the actual current state of the DC due to propagation delays). By induction, internal updates can causally depend only on internal updates, or on updates taken from the base version. Therefore, a hypothetical full version  $V_{DC} \oplus U_C$  would be causally consistent. Its projection is equivalent to the client state:  $(V_{DC} \oplus U_C)|_O = V_{DC}|_O \oplus U_C|_O = V_C$ .

This approach ensures partial availability. If a version is in the cache, it is guaranteed causally consistent, although possibly slightly stale. If it misses in the cache, the DC returns a consistent version immediately. Furthermore, the client replica can *write fast*, because it commits updates without waiting, and transfers them to its DC in the background.

Convergence is ensured, because the client's base version and log are synchronised with the DC in the background.

### 3.3 Failover and causal dependencies

The approach described so far assumes that a client connects to a single DC. However, a client can switch to a new DC at any time, in particular in response to a failure. Although each DC's state is consistent, an update that is delivered to one is not necessarily delivered in the other (because geo-replication is asynchronous, to ensure availability and performance at the DC level), which could potentially create a causal gap in the client.

To illustrate the problem, return to the example of Figure 2(a). Consider two DCs:  $DC_1$  is in (consistent) state  $V_1$ , and  $DC_2$  in (consistent) state  $V_2$ ;  $DC_1$  does not include two recent updates of  $V_2$ . Client  $C$ , connected to  $DC_2$ , replicates object  $x$  only; its state is  $V_2|_{\{x\}}$ . Suppose that the client reads the Set  $x : \{1, 3\}$ , and performs update  $u = \text{add}(4)$ , transitioning to the configuration shown in Figure 2(b).

This configuration is risky, because if the client now fails over to  $DC_1$ , and the two DCs cannot communicate, the system is not live:

<sup>6</sup>Where appropriate, we use simply  $DC$  and  $C$  to refer to some DC and client, respectively.

- (1) *Reads are not available*:  $DC_1$  cannot satisfy a request for  $y$ , since the version read by the client is more recent than the  $DC_1$  version,  $V_2 \not\subseteq V_1$ .
- (2) *Updates cannot be delivered (divergence)*:  $DC_1$  cannot deliver  $u$ , due to a missing dependency:  $u.deps \not\subseteq V_1$ .

Therefore,  $DC_1$  must reject the client (i.e., withhold its requests) to avoid creating the gap in state  $V_1 \oplus U_C$ .<sup>7</sup>

### 3.3.1 Conservative read: possibly stale, but safe

SwiftCloud provides a solution to this problem. To avoid such gaps that cannot be satisfied, the insight is to depend only on *K-stable updates* that are likely to be present in the failover DC, similarly to Mahajan et al. [30].

A version  $V$  is *K-stable* if every one of its updates is replicated in at least  $K$  DCs, i.e.,  $|\{i \in \mathcal{DC} \mid V \subseteq V_i\}| \geq K$ , where  $K \geq 1$  is a threshold configured w.r.t. expected failure model, and  $\mathcal{DC}$  is a set of all data centers. To this effect, our system maintains a consistent *K-stable version*  $V_i^K \subseteq V_i$ , which contains the updates for which  $DC_i$  has received acknowledgements from at least  $K - 1$  distinct other DCs.

A client’s base version must be *K-stable*, i.e.,  $V_C = V_{DC}^K|_O \oplus U_C|_O$ , to support failover. In this way, the client depends, either on external updates that are likely to be found in any DC ( $V_{DC}^K$ ), or internal ones, which the client can always transfer to the new DC ( $U_C$ ).

To illustrate, let us return to Figure 2(a), and consider the conservative progression to Figure 2(c), assuming  $K = 2$ . The client’s read of  $x$  returns the 2-stable version  $\{1\}$ , avoiding the dangerous dependency via an update on  $y$ . If  $DC_2$  is unavailable, the client can fail over to  $DC_1$ , reading  $y$  and propagating its update remain both live.

By the same arguments as in Section 3.2, a DC version  $V_{DC}^K$  is causally consistent and monotonically non-decreasing, and hence the client’s version as well. Note that a client observes its internal updates immediately, even if not *K-stable*.

Parameter  $K$  can be adjusted dynamically without impacting correctness. Decreasing it has immediate effect. Increasing  $K$  has effect only for future updates, to preserve monotonicity.

### 3.3.2 Discussion

The source of the above problem is an indirect causal dependency on an update to the object that the client replica does not replicate ( $y.add(2)$  in our example). As this is an inherent issue, we conjecture a general impossibility result, stating that genuine partial replication, causal consistency, partial availability and timely at-least-once delivery (convergence) are incompatible. Accordingly, the requirements must be relaxed.

Note that in many previous systems, this impossibility translates to a trade-off between consistency and availability on the one hand, and performance on the other [17, 27, 36] By “reading in the past,” we displace this to a trade-off between freshness and availability, controlled by adjusting  $K$ . A higher  $K$  increases availability, but updates take longer to be delivered;<sup>8</sup> in the limit,  $K = N$  ensures complete availability, but no client can deliver a new update when some DC is unavailable. A lower  $K$  improves freshness, but increases the probability that a client will not be able to fail over, and that it will block until its original DC recovers. In the limit,  $K = 1$

<sup>7</sup>The DC may accept to *store* client’s updates to improve durability, but it cannot deliver them or offer cache updates to the client.

<sup>8</sup>The increased number of concurrent updates that this causes is not a big problem, thanks to confluent types.

is identical to the basic protocol from Section 3.2, and is similar to blocking session-guarantee protocols [37].

$K = 2$  is a good compromise for deployments with three or more DCs, as it covers common scenarios of a DC failure or disconnection [17, 24]. Our experimental evaluation shows that it incurs a negligible staleness.

**Network partitions** Client failover between DCs is safe and generally live, except when the original set of  $K$  DCs were partitioned away from both other DCs and the client, shortly after they delivered a version to the client. In this case, the client blocks. To side-step this unavoidable possibility, we provide an unsafe API to read inconsistent data.

When a set of fewer than  $K$  DCs is partitioned from other DCs, the clients that connect to them do not deliver their mutual updates until the partition heals. To improve liveness in this scenario, SwiftCloud supports two heuristics: (i) a partitioned DC announces its “isolated” status, automatically recommending clients to use another DC, and (ii) clients who cannot reach another DC that satisfies their dependencies can use the isolated DCs with  $K$  temporarily lowered, risking unavailability if another DC fails.

**Precision vs. missing dependencies.** The probability of a client blocked due to an unsatisfied indirect causal dependency depends on many factors, such as workload- and deployment-specific ones. Representation of dependencies also contributes. SwiftCloud uses coarse-grained representation of dependencies, at the granularity of the complete base version used by the client. This may cause a *spurious* missing dependency, when a DC rejects a client because it misses some update that is not an actual dependence. Finer-grained dependency representation, such as in causality graphs [27], or resorting to application-provided explicit dependencies under a weaker variant of causal consistency [6], avoid some spurious dependencies at the expense of fatter metadata. However, the missing dependency issue remains under any dependency precision. Thus, our approach is to fundamentally minimize the chances of *any* missing dependency, both genuine and spurious.

## 4 Implementation

We now describe a metadata and concrete protocols implementing the abstract design.

### 4.1 Timestamps, vectors and log merge

The SwiftCloud approach requires metadata: (1) to *uniquely identify an update*; (2) to *encode its causal dependencies*; (3) to *identify and compare versions*; (4) and to *identify all the updates of a transaction*. We now describe a metadata design that fulfils the requirements and has a low cost. It combines the strengths of the two approaches outlined in Section 2.3, and is both *lean and safe*.

A timestamp is a pair  $(j, k) \in (\mathcal{DC} \cup \mathcal{C}) \times \mathbb{N}$ , where  $j$  identifies the node that assigned the timestamp (either a DC or a client) and  $k$  is a sequence number. Similarly to the solution of Ladin et al. [25], the metadata assigned to some update  $u$  combines both: (i) a single *client-assigned timestamp*  $u.t_C$  that uniquely identifies the update, and (ii) a set of zero or more *DC-assigned timestamps*  $u.T_{DC}$ . Before being delivered to a DC, the update has only a client-assigned timestamp, but no DC timestamp; after delivery, it has both; in case of DC failover, it may have several DC timestamps. Nodes can refer to an update via any of its timestamps in

order to tolerate failures. The updates in a transaction all have the same timestamp(s), to ensure all-or-nothing delivery [36].

We represent a version or a dependency as a *version vector*. A vector is a partial map from node ID to integer, e.g.,  $VV = [DC_1 \mapsto 1, DC_2 \mapsto 2]$ , which we interpret as a set of timestamps. For example, when  $VV$  is used as a dependency for some update  $u$ , it means that  $u$  causally depends on  $\{(DC_1, 1), (DC_2, 1), (DC_2, 2)\}$ . In SwiftCloud protocols, every vector has at most one client entry, and multiple DC entries; thus, its size is bounded by the number of DCs, limiting network overhead. In contrast to a dependency graph, a vector compactly represents transitive dependencies and can be evaluated locally by any node.

Formally, the timestamps represented by a vector  $VV$  are given by a function  $\mathcal{T}$ :

$$\mathcal{T}(VV) = \{(j, k) \in \text{dom}(VV) \times \mathbb{N} \mid k \leq VV(j)\}$$

Similarly, the version decoding function  $\mathcal{V}$  of vector  $VV$  on a state  $U$  selects every update in  $U$  that matches the vector ( $\mathcal{V}$  is defined for states  $U$  that cover all timestamps of  $VV$ ):

$$\mathcal{V}(VV, U) = \{u \in U \mid (u.T_{DC} \cup \{u.t_C\}) \cap \mathcal{T}(VV) \neq \emptyset\}$$

For the purpose of the decoding function  $\mathcal{V}$ , a given update can be flexibly referred through any of its timestamps. Moreover,  $\mathcal{V}$  is stable with growing state  $U$ . This is useful to identify a version on a large state that undergoes concurrent log appends; formally,  $\forall VV, U, U' : U \subset U' \wedge \mathcal{T}(VV) \subseteq \bigcup_{u \in U} (u.T_{DC} \cup \{u.t_C\}) \implies \mathcal{V}(VV, U) = \mathcal{V}(VV, U')$ .

The log merge operator  $U_1 \oplus U_2$ , which eliminates duplicates, is implemented using client timestamps. Two updates  $u_1 \in U_1, u_2 \in U_2$  are identical if  $u_1.t_C = u_2.t_C$ . The merge operator merges their DC timestamps into  $u \in U_1 \oplus U_2$ , such that  $u.T_{DC} = u_1.T_{DC} \cup u_2.T_{DC}$ .

## 4.2 Protocols

We now describe the protocols of SwiftCloud by following the lifetime of an update, and with reference to the names in Figure 1.

**State** A DC replica maintains its state  $U_{DC}$  in durable storage. The state respects causality and atomicity for each individual object, but due to internal concurrency, this may not be true across objects. Therefore, the DC also has a vector  $VV_{DC}$  that identifies a safe, monotonically non-decreasing causal version in the local state, which we note  $V_{DC} = \mathcal{V}(VV_{DC}, U_{DC})$ . Initially,  $U_{DC}$  contains no updates, and vector  $VV_{DC}$  is zeroed.

A client replica stores the commit log of its own updates  $U_C$ , and the projection of the base version from the DC, restricted to its interest set  $O$ ,  $V_{DC}|_O$ , as described previously in Section 3.2. It also stores a copy of vector  $VV_{DC}$  that describes the base version.

**Client-side execution** When the application starts a transaction  $\tau$  at client  $C$ , the client initialises it with an empty buffer of updates  $\tau.U = \emptyset$  and a *snapshot vector* of the current base version  $\tau.depsVV = VV_{DC}$ ; the DC can update the client's base version concurrently with the transaction execution. A read in transaction  $\tau$  is answered from the version identified by the snapshot vector, merged with recent internal updates,  $\tau.V = \mathcal{V}(\tau.depsVV, V_{DC}|_O) \oplus U_C|_O \oplus \tau.U$ . If the requested object is not in the client's interest set,  $x \notin O$ , the client extends its interest set, and returns the value once the DC updates the base version projection.

When the application issues internal update  $u$ , it is appended to the transaction buffer  $\tau.U \leftarrow \tau.U \oplus \{u\}$ , and included in any later read. To simplify the notation, we assume hereafter that

a transaction performs at most one update.<sup>9</sup> The transaction commits locally at the client and never fails. If the transaction made update  $u \in \tau.U$ , the client replica commits it locally as follows: (1) Assign it client timestamp  $u.t_C = (C, k)$ , where  $k$  counts the number of updates at the client; (2) assign it a *dependency vector* initialised with the transaction snapshot vector  $u.depsVV = \tau.depsVV$ ; (3) append it to the commit log of local updates on stable storage  $U_C \leftarrow U_C \oplus \{u\}$ . This terminates the transaction; the client can start a new one, which will observe the committed updates.

**Transfer protocol: Client to DC** The transfer protocol transmits committed updates from a client to its current DC, in the background. It repeatedly picks the first unacknowledged committed update  $u$  from the log. If any of  $u$ 's internal dependencies has recently been assigned a DC timestamp, it merges this timestamp into the dependency vector. Then, the client sends a copy of  $u$  to its current DC. The client expects to receive an acknowledgement from the DC, containing the timestamp(s)  $T$  that the DC assigned to update  $u$ . If so, the client records the timestamps in the original update record  $u.T_{DC} \leftarrow T$ . In the failure-free case,  $T$  is a singleton.

The client may now transfer the next update in the log.

A transfer request may fail for three reasons:

- (a) Timeout: the DC is suspected unavailable; the client connects to another DC (failover) and repeats the protocol.
- (b) The DC reports a *missing internal dependency*, i.e., it has not received some update of the client, as a result of a prior failover. The client recovers by marking as unacknowledged all internal updates starting from the oldest missing dependency, and restarting the transfer protocol from that point.
- (c) The DC reports a *missing external dependency*; this is also an effect of failover. In this case, the client tries yet another DC. The approach from Section 3.3.1 avoids this failing indefinitely.

Upon receiving update  $u$ , the DC verifies if its dependencies are satisfied, i.e., if  $\mathcal{T}(u.depsVV) \subseteq \mathcal{T}(VV_{DC})$ . (If this check fails, it reports an error to the client, indicating either case (b) or (c)). If the DC has not received this update previously, as determined by its client timestamp, i.e.,  $\forall u' \in U_{DC} : u'.t_C \neq u.t_C$ , the DC does the following: (1) Assign it a DC timestamp  $u.T_{DC} \leftarrow \{(DC, VV_{DC}(DC) + 1)\}$ , (2) store it in its durable state  $U_{DC} \oplus \{u\}$ , (3) make the update visible in the DC version  $V_{DC}$ , by incorporating its timestamp(s) into  $VV_{DC}$ . This last step makes  $u$  available to the geo-replication and notification protocols, described hereafter. If the update has been received before, the DC looks up its previously-assigned DC timestamps. In either case, the DC acknowledges the transfer to the client with the DC timestamp(s). Note that some of these steps can be parallelised between transfer requests received from different client replicas, e.g., using batched timestamp assignment.

**Geo-replication protocol: DC to DC** The geo-replication protocol relies on a uniform reliable broadcast across DCs. An update enters the geo-replication protocol when a DC accepts a fresh update during the transfer protocol. The accepting DC broadcasts it to all other DCs. The broadcast implementation stores an update in a replication log until every DC receives it. A DC that receives a broadcast message containing  $u$  does the following: (1) If the dependencies of  $u$  are not met, i.e., if  $\mathcal{T}(u.depsVV) \not\subseteq \mathcal{T}(VV_{DC})$ , buffer it until they are; and (2) incorporate  $u$  into durable state  $U_{DC} \oplus \{u\}$  (if  $u$  is not fresh, the duplicate-resilient log merge safely unions all timestamps), and incorporate its timestamp(s) into the DC version vector  $VV_{DC}$ . This last

<sup>9</sup>This can be easily extended to multiple updates, by assigning the same timestamp to all the updates of the same transaction, ensuring the all-or-nothing property [36].



step makes it available to the notification protocol. The  $K$ -stable version  $V_{DC}^K$  is maintained similarly.

**Notification protocol: DC to Client** A DC maintains a best-effort notification session, over a FIFO channel, to each of its connected clients. The soft state of a session includes a copy of the client’s interest set  $O$  and the last known base version vector used by the client,  $VV'_{DC}$ . The DC accepts a new session only if its own state is consistent with the base version provided by the client, i.e., if  $\mathcal{T}(VV'_{DC}) \subseteq \mathcal{T}(VV_{DC})$ . Otherwise, the client is redirected to another DC, since the DC would cause a causal gap with the client’s state (the solution from Section 3.3.1 avoids repeated rejections).

The DC sends over each channel a causal stream of update notifications.<sup>10</sup> Notifications are batched according to either time or to rate [9]. A notification packet consists of a new base version vector  $VV_{DC}$ , and a log of all the updates  $U_\delta$  to the objects of the interest set, between the client’s previous base vector  $VV'_{DC}$  and the new one. Formally,  $U_\delta = \{\mathbf{u} \in U_{DC|O} \mid \mathbf{u}.T_{DC} \cap (\mathcal{T}(VV_{DC}) \setminus \mathcal{T}(VV'_{DC})) \neq \emptyset\}$ . The client applies the newly-received updates to its local state described by the old base version,  $V_{DC|O} \leftarrow V_{DC|O} \oplus U_\delta$ , and assumes the new vector  $VV_{DC}$ . If any of received updates is a duplicate w.r.t. to the old version or to a local update, the log merge operator handles it safely. Note that transaction atomicity is preserved, since all updates of a transaction share a common timestamp, thus either all fit in a batch  $U_\delta$  or none does.

When the client detects a broken channel, it reinitiates the session, possibly on a new DC.

The interest set can change dynamically. When an object is evicted from the cache, the notifications are lazily unsubscribed to save resources. When it is extended with object  $x$ , the DC responds with the current version of  $x$ , which includes all updates to  $x$  up to the base version vector. To avoid races, a notification includes a hash of the interest set, which the client checks.

### 4.3 Object checkpoints and log pruning

Update logs contribute to substantial storage and, to smaller extent, network costs. To avoid unbounded growth, a *pruning protocol* periodically replaces the prefix of a log by a *checkpoint*. In the common case, a checkpoint is more compact than the corresponding log of updates; for instance, a log containing one thousand increments to a Counter object, each with its timestamps, can be replaced by a checkpoint containing just the number 1000 paired with a version vector.

#### 4.3.1 Log pruning in the DC

The log at a DC provides (a) protection from duplicated update delivery, as abstracted by the  $\oplus$  operator, and (b) the capability to compute different versions, for application processes reading at different causal times. A log entry for update  $\mathbf{u}$  may be replaced with a checkpoint, once all of its duplicates have been filtered out, and once  $\mathbf{u}$  has been delivered to all interested application processes.

Precise evaluation of this condition would require access to the client replica states. In practice, we need to *prune aggressively*, but without violating correctness. In order to reduce the risk of pruning a version not yet delivered to an interested application (which could force it to restart an ongoing transaction), we prune only a *delayed version*  $VV_{DC}^\Delta$ , where  $\Delta$  is a real-time delay [27, 28]. If this heuristic fails, the consequences are not fatal: an ongoing client

<sup>10</sup>Alternatively, the client can ask for invalidations instead, trading responsiveness for lower bandwidth utilization and higher DC throughput.

	<b>YCSB</b> [16]	<b>SocialApp</b> [36]
Type of objects	LWW Map	Set, Counter, Register
Object payload	$10 \times 100$ bytes	variable
Read transactions	read fields (workload A: 50% / B: 95%)	read wall <sup>†</sup> (80%) see friends (8%)
Update transactions	update field (workload A: 50% / B: 5%)	message (5%) post status (5%) add friend (2%)
Objects / transaction	1 (non-transactional)	2–5
Database size	50,000 objects	50,000 users (400,000 objects)
Object popularity	uniform / Zipfian	uniform
Session locality	40% (low) / 80% (high)	

<sup>†</sup> *Read wall* is an update if page view statistics are enabled.

Table 1: Characteristics of applications/workloads.

transaction may need to restart and repeat prior reads or return inconsistent data if desired, but the committed updates are *never* aborted.

To avoid duplicates despite log pruning, we extend DC metadata as follows.  $DC_i$  maintains an *at-most-once guard*  $G_i : \mathcal{C} \rightarrow \mathbb{N}$ , which records the sequence number of each client’s last pruned update. The guard is local to and shared at a DC. Whenever the DC receives a transfer request or a geo-replication message for update  $u$  with client timestamp  $(C, k)$  and cannot find it in its log, it checks the at-most-once guard  $G_i(C)$  entry. If the DC recognises that the update is a duplicate of a pruned update ( $G_i(C) \geq k$ ), it ignores the update, except that it advances version vector to include all of  $u$ ’s DC timestamps; for a transfer request, the DC replies with a vector  $VV_i$ , which is an overapproximation of the (discarded) information about the exact set of  $u$ ’s DC timestamps.

The notification protocol also uses checkpoints. On a client cache miss, instead of a complete log, the DC sends an equivalent checkpoint of the object, together with the client’s guard entry, so that the client can merge it with its log safely.

### 4.3.2 Pruning the client’s log

Managing the log at a client is simpler. A client logs its own updates  $U_C$ , which may include updates to object that is currently out of its interest set. This enables the client to read its own updates, and to propagate them lazily to a DC when connected and convenient. An update  $u$  can be discarded as soon as it appears in the  $K$ -stable base version  $V_{DC}^K$ , i.e., when the client becomes dependent on the presence of  $u$  at a DC. The client discards the corresponding updates:  $U_C \leftarrow U_C \setminus \mathcal{V}(VV_i^K, U_C)$ .

## 5 Evaluation

We implement SwiftCloud and evaluate it experimentally, in comparison to alternatives. We show that SwiftCloud provides: (i) fast response, under 1 ms for both reads and writes to cached objects (Section 5.3); (ii) while supporting thousands of clients, throughput that scales with the number of DCs and small metadata, sized linear in the number of DCs (Section 5.4);

(iii) fault-tolerance w.r.t. client churn (Section 5.5) and DC outages (Section 5.6); and (iv) low staleness, under 1% of stale reads under common conditions (Section 5.7).

## 5.1 Implementation and applications

SwiftCloud and the benchmark applications are implemented in Java.<sup>11</sup> SwiftCloud uses an extendable library of log-based CRDT types [35], in-memory storage,<sup>12</sup> Kryo for data marshalling, and a custom RPC implementation. A client cache has a fixed size and uses an LRU eviction policy. More elaborate approaches, such as object prefetching [9], are feasible.

Our client API resembles both production object stores, such as Riak 2.0 or Redis [1, 32], and prototype causal transactional stores, such as COPS or Eiger [27, 28]:<sup>13</sup>

<code>begin_transaction()</code>	<code>read(object) : value</code>
<code>commit_transaction()</code>	<code>update(object, method(args...))</code>

The actual API also includes caching options omitted here.

Along the lines of previous studies of causally-consistent systems [4, 6, 28, 36], we use two different benchmarks, YCSB and SocialApp, summarised in Table 1.

YCSB [16] serves as a kind of micro-benchmark, with simple requirements, measuring baseline costs and specific system properties in isolation. It has a simple key-field-value object model, which we implement as a LWW Map type, using a default payload of ten fields of 100 bytes each. YCSB issues single-object reads and writes. We use two of the standard YCSB workloads: update-heavy Workload A, and read-dominated Workload B. The object access pattern can be set to either uniform or Zipfian. YCSB does not rely on transactional semantics or high-level data types.

SocialApp is a social network application modelled after WaltSocial [36].<sup>14</sup> It employs high-level data types such as Sets, for friends and posts, LWW Register for profile information, Counter for counting profile visits, and object references. Many SocialApp objects grow in size over time (e.g., sets of posts). We are not concerned about this growth; a recent work of Briquemont [11] demonstrates how to implement object sharding in a SwiftCloud-like system. SocialApp accesses multiple objects in a causal transaction to ensure that operations such as reading a wall page and profile information behave consistently. Percentage in parentheses indicate the frequency of each operation in the workload. The SocialApp workload is read-dominated, but note that visiting a wall actually increments the wall visit counter when statistics are enabled, in metadata experiments. The user popularity distribution is uniform.

We use a 50,000-user database for both applications, except for smaller 10,000 users database for metadata experiments, to increase the stability of measurements.

We are not aware of any realistic benchmark for large-scale client-side apps that would define a workload with thousands of long client sessions. We evaluate the system with short client sessions issuing more frequent operations than we expect in realistic workloads. The system behaviour under such a condensed workload is our proxy for the behaviour with more clients running slower sessions of longer duration.

In order to model the locality behaviour of a client, both YCSB and SocialApp are augmented with a facility to control access locality, mimicking social network access patterns [10]. Within a

<sup>11</sup> <https://github.com/SyncFree/SwiftCloud>

<sup>12</sup> Our prototype can use BerkeleyDB for durable storage, but it was turned off in the present experiments.

<sup>13</sup> Unlike COPS or Eiger, we consider interactive transactions, i.e., accessed objects do not need to be predefined.

<sup>14</sup> SocialApp does not implement WaltSocial’s user registration operation, as it would require additional support for strong consistency.

client session, a workload generator draws uniformly from a pool of session-specific objects with either 40% (*low locality*) or 80% (*high locality*) probability. For SocialApp, the pool contains objects of user’s friends. Objects not drawn from this local pool are drawn from the global distribution (uniform or Zipfian) described above. The local pool can fit in the client’s cache.

## 5.2 Experimental setup

We run three DCs in geographically distributed Amazon EC2 availability zones (Europe, Virginia, and Oregon), and a pool of distributed clients. Round-Trip Times (RTTs) between nodes are as follows:

	Oregon DC	Virginia DC	Europe DC
nearby clients	60–80 ms	60–80 ms	60–80 ms
Europe DC	177 ms	80 ms	
Virginia DC	60 ms		

Each DC runs on a single m3.m EC2 instance, cheap virtual hardware, equivalent to a single core 64-bit 2.0 GHz Intel Xeon processor (2 ECUs) with 3.75 GB of RAM, and OpenJDK7 on Linux 3.2. Objects are pruned at random intervals between 60–120 s, to avoid bursts of pruning activity. We deploy 500–2,500 clients on a separate pool of 90 m3.m EC2 instances. Clients load DCs uniformly and use the closest DC by default, with a client-DC RTT ranging in 60–80 ms.

For comparison, we provide three protocol modes: (i) *SwiftCloud mode* (default) with client cache replicas of 256 objects, refreshed with notifications at a rate  $\leq 1$  s by default; (ii) *Safe But Fat metadata mode* with cache, but with client-assigned metadata only (modelled after PRACTI, or Depot without cryptography [9, 30]), (iii) *server-side replication mode* without client caches; in this mode, a read incurs one RTTs to a DC, whereas an update incurs two RTTs to a DC, modelling the cost of a synchronous write to a quorum of servers to ensure fault-tolerance comparable to SwiftCloud.

## 5.3 Response time and throughput

We run several experiments to compare SwiftCloud’s client-side caching, with reference to the locality potential and server-side geo-replication without caching. For each workload we evaluate the system stimulated with different rates of aggregated incoming transactions, until it becomes saturated. We use a number of clients that is throughput-optimised for each pair of workload and protocol mode. We report aggregated statistics for all clients.

Figure 3 shows response times for YCSB, comparing server-only (left side) with client replication (right side), under low (top) and high locality (bottom), when the system is not overloaded. Recall that in server-only replication, a read incurs a RTT to the DC, whereas an update incurs 2 RTTs. We expect SwiftCloud to provide much faster response, at least for cached data. Indeed, the figure shows that a significant fraction of operations respond immediately in SwiftCloud mode, and this fraction tracks the locality of the workload (marked “locality potential” on the figure), within a  $\pm 7.5$  percentage-point margin attributable to caching policy artefacts.<sup>15</sup> The remaining operations require one round-trip to the DC, indicated as 1 RTT. As our measurements for SocialApp show the same message, we do not report them here. These results demonstrate that the consistency guarantees and the rich programming interface of SwiftCloud do *not* affect responsiveness of read and update of cached data.

<sup>15</sup>A detailed analysis reveals the sources of this error margin. The default Zipfian object access distribution of YCSB increases the fraction of local accesses due to added “global” locality (up to 82% local accesses for target




Figure 3: Response time for YCSB operations (workload A, Zipfian object popularity) under different system and workload locality configurations, aggregated for all clients.




Figure 4: Maximum system throughput for different workloads and protocols. Percentage over bars indicates the increase/decrease in throughput for SwiftCloud compared to server-side replication.




Figure 5: Throughput vs. response time for different system configurations running variants of YCSB.

In the next study, we saturate the system to determine its maximum aggregated throughput. Figure 4 compares SwiftCloud with server-side replication for all workloads.

Client-side replication is a mixed blessing: client replicas absorb read requests that would otherwise reach the DC, but on the other hand require additional work at the DC to maintain client replicas. The cost of client replicas pays off for read-dominated high locality workloads. SwiftCloud consistently delivers higher throughput for high locality workloads, by 7% up to 128%, and for read-heavy workloads in particular. In contrast, low locality workloads show no clear trend; depending on the workload, throughput either increases by up to 38%, or decreases by up to 11%.

Our next experiment studies how response times vary with server load and with the staleness settings. The results show that, as expected, cached objects respond immediately and are always available, but the responsiveness of cache misses depends on server load. For this study, Figure 5 plots throughput vs. response time, for YCSB A (left side) and B (right side), both for the Zipfian (top) and uniform (bottom) distributions. Each point represents the aggregated throughput and latency for a given transaction incoming rate, which we increase until reaching the saturation point. The curves report two percentiles of response time: the lower (70th percentile) line represents the response time for requests that hit in the cache (the session locality level is 80%), whereas the higher (95th percentile) line represents misses, i.e., requests served by a DC.

As expected, the lower (cached) percentile consistently outperforms the server-side baseline, for all workloads and transaction rates. A separate analysis, not reported in detail here, reveals that a saturated DC slows down its rate of notifications, increasing staleness, but this does not impact response time, as desired. In contrast, the higher percentile follows the trend of server-side replication response time, increasing remote access time.

80% of workload session locality). On the other hand, low locality workload decreases amount of local accesses, due to magnified imperfections of LRU cache eviction algorithm (down to 34,% local accesses for target 40%).




Figure 6: Maximum system throughput for a variable number of client and DC replicas.

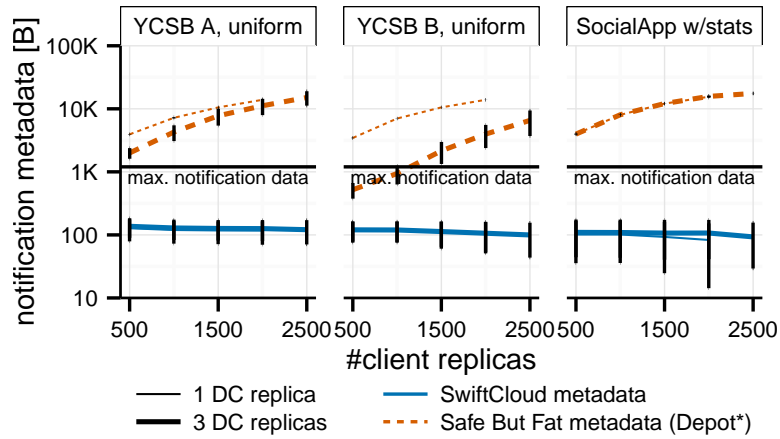


Figure 7: Size of metadata in notification message for a variable number of replicas, mean and standard error. Normalised to a notification of 10 updates.

Varying the target notification rate (not plotted) between 500 ms and 1000 ms, reveals the same trend: response time is not affected by the increased staleness. At a lower refresh rate, notification batches are less frequent but larger. This increases throughput for the update-heavy YCSB A (up to tens of percent points), but has no effect on the throughput of read-heavy YCSB B. We expect the impact of refresh rate to be amplified for workloads with smaller rate of notification updates.

## 5.4 Scalability

Next, we measure how well SwiftCloud scales with increasing numbers of DC and of client replicas. Of course, performance is expected to increase with more DCs, but most importantly, the size of metadata should be small, should not depend on the number of clients, and should increase linearly with the number of DCs. Our results confirm these expectations.

In this experiment, we run SwiftCloud with a variable number of client (500–2500) and server (1–3) replicas. We report only on the uniform object distribution, because under the

Zipfian distribution different numbers of clients skew the load differently, making any comparison meaningless. To control staleness, we run SwiftCloud with two different notification rates (every 1 s and every 10 s).

Figure 6 shows the maximum system throughput on the Y axis, increasing the number of replicas along the X axis. The thin lines are for a single DC, the bold ones for three DCs. Solid lines represent the fast notification rate, dashed lines the slow one. The figure shows, left to right, YCSB Workload A, YCSB Workload B, and SocialApp.

The capacity of a single DC in our hardware configuration peaks at 2,000 active client replicas for YCSB, and 2,500 for SocialApp. Beyond that, the DC drops sessions.

The wisdom of server-side replication applies in a new way to SwiftCloud: additional DC replicas increase the system capacity for operations that can be performed at only one replica. These traditionally include only read operations, but in the case of SwiftCloud, also sending notification messages (maintaining active clients). Whereas a single SwiftCloud DC supports at most 2,000 clients. With three DCs SwiftCloud supports at least 2,500 clients for all workloads. Unfortunately, as we ran out of resources for client machines at this point, we cannot report an upper bound.

For some fixed number of DCs, adding client replicas increases the aggregated system throughput, until a point of approximately 300–500 clients per DC, where the cost of maintaining client replicas up to date saturates the DCs, and further clients do not absorb enough reads to overcome that cost. Note that the lower refresh rate can control the load at a DC, and reduces it by 5 to 15%.

In the same experiment, Figure 7 presents the distribution of metadata size in notification messages. (Notifications are the most common and the most costly messages sent over the network.) We plot the size of metadata (in bytes) on the Y axis, varying the number of clients along the X axis. Left to right, the same workloads as in the previous figure. Thin lines are for one DC, thick lines for three DCs. A solid line represents SwiftCloud “Lean and Safe” metadata, and dotted lines the classical “Safe But Fat” approach. Note that our Safe-but-Fat implementation includes the optimisation of sending modified entries of vector rather than the full vector, as in Depot or PRACTI [9, 30]. Vertical bars represent standard error across clients. As notifications are batched, we normalise metadata size to a message carrying exactly 10 updates, corresponding to under approx. 1 KB of data.

This plot confirms that the SwiftCloud metadata is small and constant, at 100–150 bytes/notification (10–15 bytes per update); data plus metadata together fit inside a single standard network packet.<sup>16</sup> It is *independent* both from the number of client replicas and from the workload, as well as from the number of objects in the database, as an additional experiment (not plotted) validates. Increasing the number of DC replicas from one to three causes a negligible increase in metadata size, of under 10 bytes. We attribute some variability to the data encoding and inaccuracies of measurements, including the normalisation process.

In contrast, the classical Safe-but-Fat metadata grows linearly with the number of clients and exhibits higher variability. Its size reaches approx. 1 KB for 1,000 clients in all workloads, and 10 KB for 2,500 clients. Clearly, metadata that is up to 10× larger than the actual data represents a substantial overhead.

<sup>16</sup>The size of metadata does not exceed 100–150 bytes per notification with 10 updates, which matches our back of the envelope computations: 2–4 entries in a vector, plus 10 pairs of timestamps (one for each update) yields approximately 24 timestamps in total, with potential duplicates encoded more efficiently.






Figure 8: Storage occupation at one DC in reaction to client churn, for SwiftCloud and Lean-but-Unsafe alternative.

## 5.5 Tolerating client churn

We now turn to fault tolerance. In the next experiment, we evaluate SwiftCloud under client churn, by periodically disconnecting client replicas and replacing them with a new set of clients. At any point in time, there are 500 active clients and a variable number of disconnected clients, up to 5000. Figure 8 illustrates the storage occupation of a DC for representative workloads, which is also a proxy for the size of object checkpoints transferred. We compare SwiftCloud’s log compaction to a protocol without at-most-once delivery guarantees (Lean But Unsafe).

SwiftCloud storage size is approximately constant thanks to the aggressive log compaction. This is safe thanks to the at-most-once guard table per DC. Although the size of the guard (bottom curve) grows with the number of clients, it requires orders of magnitude less storage than the actual database itself.

A protocol without at-most-once delivery guarantees uses Lean-but-Unsafe metadata, without SwiftCloud’s at-most-once guard. This requires more complexity in each object’s implementation, to protect itself from duplicates. This increases the size of objects, impacting both storage and network costs. As is visible in the figure, the cost depends on the object type: none for YCSB’s LWW-Map, which is naturally idempotent, vs. linear in the number of clients for SocialApp’s Counter objects.

We conclude that the cost of maintaining SwiftCloud’s at-most-once guard is negligible, and easily amortised by its stable behaviour and possible savings.

## 5.6 Tolerating DC failures

The next experiment studies the behaviour of SwiftCloud when a DC disconnects. The scatter-plot in Figure 9 shows the response time of a SocialApp client application as the client switches between DCs. The client runs on a private machine outside of EC2. Each dot represents the response time of an individual transaction. Starting with a cold cache, response times quickly drops to near zero for transactions hitting in the cache, and to around 110 ms for misses. Some 33s into the experiment, the current DC disconnects, and the client is diverted to another DC in a different continent. Thanks to  $K$ -stability the failover succeeds, and the client continues with the new DC. Response time for cache misses reflects the higher RTT to the new DC. At 64s, the client switches back the initial DC, and performance smoothly recovers.

Recall that a server-side geo-replication system with similar fault-tolerance incurs high response time (cf. Section 5.3, or Reference 17) and does not ensure at-most-once delivery.


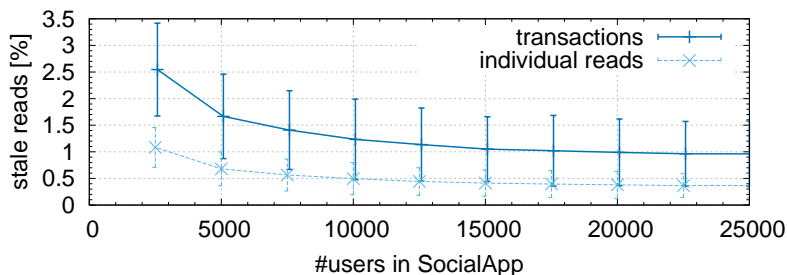


Figure 9: Response time for a client that hands over between DCs during a 30 s failure of a DC.

Figure 10:  $K$ -stability staleness overhead.

## 5.7 Staleness cost

The price to pay for our read-in-the-past approach is an increase in staleness. We consider a read *stale* if there exists, at the client's current DC, a non- $K$ -stable version that is more recent than the one returned. A transaction is stale if any of its reads is stale. In the experiments so far, we observed a negligible number of stale reads and transactions, below 1%. In another experiment, we artificially increase the probability of staleness by various means, e.g., using a smaller database, and setting cache size to zero. We run the SocialApp benchmark with 1000 clients in Europe connected to the Ireland DC and replicated in the Oregon DC.

Figure 10 shows that stale reads and stale transactions remain under 1% and 2.5% respectively. This shows that even under high contention, accessing a slightly stale snapshot has very little impact on the data read by transactions.

## 6 Related work

### 6.1 Related results on consistency and availability

Mahajan et al. [29] prove that no stronger consistency model than causal consistency is available and convergent, under *full* replication. We conjecture that these properties are not simultaneously achievable under *partial* replication, and we show how to weaken one of the liveness properties. Bailis et al. [7] also study variants of weak consistency models, and formulate a similar impossibility for a client switching server replicas. However, they do not take into account the capabilities of a client replica, as we do.

Some operations or objects of application may require stronger consistency, which requires synchronous protocols [21]. For instance, we observe that our social network application port

	System	Partial replication support	Timestamp assignment	Causality metadata size $O(\#\text{entries})$	Version metadata validity	$\leq 1$ delivery	CRDTs support
<b>Safe but</b>	PRACCTI [9]	arbitrary, no sharding	any replica	#replicas $\approx 10^4$ - $10^5$	global	yes	possibly yes
	Depot [30]	partial data at client	any replica	#replicas $\approx 10^4$ - $10^5$	global	yes	possibly yes
<b>Rat</b>	COPS-GT [27]	DC sharding	database client	causality subgraph	local (DC)	yes	possibly yes
	Bolt-on [6]	external sharding	DC server	causal dependencies	local (DC)	no	LWW only
<b>Lean but</b>	Eiger [28]	DC sharding	DC server (shard)	#objects $\approx 10^6$	local (DC)	yes <sup>†</sup>	counter, LWW
	Orbe [19]	DC sharding	DC server (shard)	#servers $\approx 10^2$ - $10^3$	global <sup>†</sup>	no	LWW only
<b>Unsafe</b>	ChainReaction [4]	DC sharding	DC (full replica)	#DCs $\approx 10^0$ - $10^1$	local (DC)	no	LWW only
	Walter [36]	arbitrary, no sharding	DC	#DCs $\approx 10^0$ - $10^1$	global	no	LWW only
<b>Lean and Safe</b>	Lazy Replication [25]	no support	DC (full replica) + client	#DCs $\approx 10^0$ - $10^1$ + 1 client entry	global	yes	single object
	SwiftCloud	no DC sharding, partial at client	DC (full replica) + client replica	#DCs $\approx 10^0$ - $10^1$ + 1 client entry	global	yes	yes

<sup>†</sup> not live during DC failures  
<sup>‡</sup> only for server-side clients

Table 2: Qualitative comparison of metadata used by the most related causally consistent systems. “Timestamp assignment” indicates which nodes assign timestamps. “Causality metadata” indicates the type and maximum cardinality of entries in causality metadata, given as the expected order of magnitude (e.g., we expect hundreds to thousands servers). “Version metadata” indicates if a metadata to represent a consistent version is valid only in a local replica (DC), or globally; for the latter, if it is fault-tolerant or not. “ $\leq 1$  delivery” indicates at-most-once delivery support.

would benefit from strongly consistent support for user registration or a password change. Prior work demonstrates that combining strong and weak consistency is possible [26, 36]. In particular, Balegas et al. [8] implemented protocols that enforce strong consistency invariants as a middleware layer on top of SwiftCloud.

## 6.2 Comparison with other systems

Several systems support consistent, available and convergent data access, at different scales. For the most related causally-consistent systems, we compare their metadata in Table 2. The columns indicate: (i) the extent of support for partial replication; (ii) which nodes assign timestamps; (iii) the worst-case size of causality metadata; (iv) the scope of validity of metadata representing a database version (is it valid only in a local replica or anywhere); (v) whether it ensures at-most-once delivery; (vi) whether it supports general confluent types (CRDTs).

### 6.2.1 Replicated databases for client-side apps

PRACTI [9] is a seminal work on causal consistency under partial replication. PRACTI uses Safe-but-Fat client-assigned metadata and a flexible log-exchange protocol that supports arbitrary communication topologies and modes. While PRACTI is very general, it is not viable for large-scale client-side replication deployment: (i) Its fat metadata approach (version vectors sized as the number of clients) is prohibitively expensive (see Figure 7), and (ii) any replica can make another unavailable, because of the indirect dependence issue discussed in Section 3.3.2. Our cloud-backed support of client-side replication addresses these issues at the cost of lower flexibility in communication topology. We are investigating support for a limited form of peer-to-peer communication that would not cause these issues, e.g., between devices of a same user or a local group of collaborators.

Our high availability techniques are similar to Depot [30], a causally-consistent storage for the client-side, built on top of untrusted cloud replicas. Depot tolerates Byzantine cloud behaviour using cryptographic metadata signatures, in order to detect misbehaviour, and fat metadata, in order to support direct client-to-client communication. Conservatively, Depot either exposes updates signed by  $K$  different servers or forces clients to receive all transitive causal dependencies of their reads. This is at odds with genuine partial replication. Under no failures, a client receives metadata of *every* update; under failures, it may also receive their body. In contrast, SwiftCloud relies on DCs to compute  $K$ -stable consistent versions with lean metadata. In the event of an extensive failure involving  $K$  DCs, SwiftCloud provides the flexibility to decrease  $K$  dynamically or to weaken consistency.

Both PRACTI and Depot systems use Safe-but-Fat metadata, as indicated in Table 2. They support only LWW registers, but their rich metadata could conceivably accommodate high-level CRDTs too.

Lazy Replication (LR) protocols [25] support multiple consistency modes for client-side apps executing operations on server replicas. Under causal consistency, LR provides high availability with asynchronous read and write requests to multiple servers. As suggested by Ladin et al. [25], LR could also read stable updates for availability on failover, but that would force its clients to execute updates synchronously. The implementation of LR uses safe and lean metadata similar to SwiftCloud, involving client- and server-assigned timestamps together with vector summaries. A log compaction protocol relies on availability of client replicas and loosely-synchronised clocks for progress.

SwiftCloud structures the database into smaller CRDT objects, which allows it to provide partial client replicas, whereas LR considers only global operations. We show that client replicas can offer higher responsiveness on cached objects, instead of directing all operations to the server side as in LR, and that local updates can be combined with  $K$ -stable updates into a consistent view, avoiding slow and unavailable synchronous updates of LR. The log compaction technique of LR is complementary to ours, and optimises for the average case. SwiftCloud’s aggressive pruning relies on at-most-once guard table, optimising for failure scenarios.

Recent web and mobile application frameworks, such as TouchDevelop [12], Google Drive Realtime API [14], or Mobius [15] support replication for in-browser or mobile applications. These systems are designed for small objects [14], database that fits on a mobile device [12], or a database of independent objects [15]. It is unknown if/how they support multiple DCs and fault tolerance, whereas SwiftCloud supports consistency, a large database, and fault tolerance. TouchDevelop provides a form of object composition, and offers integration with strong consistency [12]. We are looking into ways of adapting similar mechanisms.

### 6.2.2 Server-side geo-replicated databases

A number of geo-replicated systems offer available causally consistent data access inside a DC with excellent scale-out by sharding [4, 6, 19, 20, 27, 28].

Table 2 shows that server-side systems use variety of types of metadata, mostly Lean-but-Unsafe metadata. COPS [27] assigns metadata directly at database clients, and uses explicit dependencies (a graph). Follow-up work shows that this approach is costly, and assigns metadata at object/shard replicas instead [19, 28], or on a designated node in the DC [4, 36]. The location of assignment directly impacts the size of causality metadata. In most systems, it varies with the number of reads, with the number of dependencies, and with the stability conditions in the system. When fewer nodes assign metadata, it tends to be smaller (as in SwiftCloud), but this may limit throughput. Recent work of Du et al. [20] make use of full stability, a special case of  $K$ -stability, to remove the need for dependency metadata in messages, thereby improving throughput.

Server-side designs do not easily extend beyond the scope and the failure domain of a DC, because (i) their protocols do not tolerate external client failures and DC outages, either blocking or violating safety (due to inadequate metadata, and the causal dependence issue); (ii) as they assume that data is updated by overwriting, implementing high-level confluent data types that work on the client-side is complex and costly (see Figure 8); (iii) their metadata can grow with database size.

SwiftCloud’s support for server-side sharding is limited compared to the most scalable geo-replicated designs. Reconciling client-side replication with a more decentralised sharding support, and small metadata size, is future work. We believe this is possible to achieve by, once again, trading data freshness for performance, i.e., by managing a slightly stale consistent version at a high throughput, with small metadata [20].

## 7 Conclusion

We presented the design of SwiftCloud, the first object database that offers client-side apps a local access to partial replica with the guarantees of geo-replicated systems.

Our experiments show that the design of SwiftCloud is able to provide immediate and consistent response for reads and updates on local objects, and to maintain the throughput of a

server-side geo-replication, or better. SwiftCloud’s metadata allows it to scale safely to thousands of clients with 3 DCs, with small size objects, and metadata at the level of 15 bytes per update, independent of the number of connected and disconnected clients. Our fault-tolerant protocols handle failures nearly transparently, at a low staleness cost.

SwiftCloud’s design leverages a common principle that helps to achieve several goals: client buffering and controlled staleness can absorb the cost of scalability, availability, and consistency.

Several aspects remain open for improvement and investigation. Our DC implementation is not sharded — we wish to combine modern sharded DC protocols with SwiftCloud, ideally without increasing the size of metadata. Practical applications require security mechanisms; we expect to adapt Depot’s support for Byzantine clients, and additional access control and privacy mechanisms at the object level. We are also looking to better integration with programming models, in particular in terms of support for mixed weak and strong consistency, and for object composition [22].

**Acknowledgments** We would like to thank Carlos Baquero, Peter Bailis, Allen Clement, Alexey Gotsman, Masoud Saeida Ardekani, João Leitão, Vivien Quéma, and Luís Rodrigues for their comments that helped to improve this work. This research is supported in part by ANR (France) project ConcoRDanT (ANR-10-BLAN 0208), by Google Europe Fellowship in Distributed Computing 2010 awarded to Marek Zawirski, and by European FP7 project SyncFree (project 609 551, 2013–2016). ☐

## References

- [1] Introducing Riak 2.0: Data types, strong consistency, full-text search, and much more, Oct. 2013. URL <http://basho.com/introducing-riak-2-0/>.
- [2] M. Ahamad, J. E. Burns, P. W. Hutto, and G. Neiger. Causal memory. In *Proc. 5th Int. Workshop on Distributed Algorithms*, pages 9–30, Delphi, Greece, Oct. 1991.
- [3] P. S. Almeida and C. Baquero. Scalable eventually consistent counters over unreliable networks. Number arXiv:1307.3207, July 2013. URL <http://arxiv.org/abs/1307.3207>.
- [4] S. Almeida, J. Leitão, and L. Rodrigues. ChainReaction: a causal+ consistent datastore based on Chain Replication. In *Euro. Conf. on Comp. Sys. (EuroSys)*, Apr. 2013.
- [5] P. Bailis and K. Kingsbury. The network is reliable: An informal survey of real-world communications failures. *ACM Queue*, 2014.
- [6] P. Bailis, A. Ghodsi, J. M. Hellerstein, and I. Stoica. Bolt-on causal consistency. In *Int. Conf. on the Mgt. of Data (SIGMOD)*, pages 761–772, New York, NY, USA, 2013. doi: 10.1145/2463676.2465279.
- [7] P. Bailis, A. Davidson, A. Fekete, A. Ghodsi, J. M. Hellerstein, and I. Stoica. Highly Available Transactions: Virtues and limitations. In *Int. Conf. on Very Large Data Bases (VLDB)*, Riva del Garda, Trento, Italy, 2014.
- [8] V. Balesgas, N. Preguiça, R. Rodrigues, S. Duarte, C. Ferreira, M. Najafzadeh, and M. Shapiro. Putting the consistency back into eventual consistency. In *Euro. Conf. on Comp. Sys. (EuroSys)*, Bordeaux, France, Apr. 2015.
- [9] N. Belaramani, M. Dahlin, L. Gao, A. Nayate, A. Venkataramani, P. Yalagandula, and J. Zheng. PRACTI replication. In *Networked Sys. Design and Implem. (NSDI)*, pages 59–72, San Jose, CA, USA, May 2006. Usenix, Usenix.
- [10] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user behavior in online social networks. In *Internet Measurement Conference (IMC)*, 2009.

- 
- [11] I. Briquemont. Optimising Client-side Geo-replication with Partially Replicated Data Structures. Master's thesis, Universite Catholique de Louvain, Louvain-la-Neuve, Belgium, 2014.
- [12] S. Burckhardt. Bringing TouchDevelop to the cloud. Inside Microsoft Research Blog, Oct. 2013. URL [http://blogs.technet.com/b/inside\\_microsoft\\_research/archive/2013/10/28/bringing-touchdevelop-to-the-cloud.aspx](http://blogs.technet.com/b/inside_microsoft_research/archive/2013/10/28/bringing-touchdevelop-to-the-cloud.aspx).
- [13] S. Burckhardt, A. Gotsman, H. Yang, and M. Zawirski. Replicated data types: Specification, verification, optimality. In *Symp. on Principles of Prog. Lang. (POPL)*, pages 271–284, San Diego, CA, USA, Jan. 2014. doi: 10.1145/2535838.2535848.
- [14] B. Cairns. Build collaborative apps with Google Drive Realtime API. Google Apps Developers Blog, Mar. 2013. URL <http://googleappsdeveloper.blogspot.com/2013/03/build-collaborative-apps-with-google.html>.
- [15] B.-G. Chun, C. Curino, R. Sears, A. Shraer, S. Madden, and R. Ramakrishnan. Mobius: Unified messaging and data serving for mobile apps. In *Int. Conf. on Mobile Sys., Apps. and Services (MobiSys)*, pages 141–154, New York, NY, USA, 2012.
- [16] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with YCSB. In *Symp. on Cloud Computing*, pages 143–154, Indianapolis, IN, USA, 2010.
- [17] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, W. Hsieh, S. Kanthak, E. Kogan, H. Li, A. Lloyd, S. Melnik, D. Mwaura, D. Nagle, S. Quinlan, R. Rao, L. Rolig, Y. Saito, M. Szymaniak, C. Taylor, R. Wang, and D. Woodford. Spanner: Google's globally-distributed database. In *Symp. on Op. Sys. Design and Implementation (OSDI)*, pages 251–264, Hollywood, CA, USA, Oct. 2012. Usenix.
- [18] 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 *Symp. on Op. Sys. Principles (SOSP)*, volume 41 of *Operating Systems Review*, pages 205–220, Stevenson, Washington, USA, Oct. 2007. Assoc. for Computing Machinery. doi: <http://doi.acm.org/10.1145/1294261.1294281>.
- [19] J. Du, S. Elnikety, A. Roy, and W. Zwaenepoel. Orbe: Scalable causal consistency using dependency matrices and physical clocks. In *Symp. on Cloud Computing*, pages 11:1–11:14, Santa Clara, CA, USA, Oct. 2013. Assoc. for Computing Machinery. doi: 10.1145/2523616.2523628.
- [20] J. Du, C. Iorgulescu, A. Roy, and W. Zwaenepoel. GentleRain: Cheap and scalable causal consistency with physical clocks. In *Symp. on Cloud Computing*, 2014.
- [21] S. Gilbert and N. Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, 2002. ISSN 0163-5700. doi: <http://doi.acm.org/10.1145/564585.564601>.
- [22] A. Gotsman and H. Yang. Composite replicated data types. In *Euro. Symp. on Programming (ESOP)*, London, UK, 2015.
- [23] P. R. Johnson and R. H. Thomas. The maintenance of duplicate databases. Internet Request for Comments RFC 677, Information Sciences Institute, Jan. 1976. URL <http://www.rfc-editor.org/rfc.html>.
- [24] A. Kansal, B. Urgaonkar, and S. Govindan. Using dark fiber to displace diesel generators. In *Hot Topics in Operating Systems*, Santa Ana Pueblo, NM, USA, 2013.
- [25] R. Ladin, B. Liskov, and L. Shrira. Lazy replication: Exploiting the semantics of distributed services. *Operating Systems Review*, 25(1):49–55, Jan. 1991.

- [26] C. Li, D. Porto, A. Clement, J. Gehrke, N. Preguiça, and R. Rodrigues. Making geo-replicated systems fast as possible, consistent when necessary. In *Symp. on Op. Sys. Design and Implementation (OSDI)*, pages 265–278, Hollywood, CA, USA, Oct. 2012.
- [27] W. Lloyd, M. J. Freedman, M. Kaminsky, and D. G. Andersen. Don't settle for eventual: scalable causal consistency for wide-area storage with COPS. In *Symp. on Op. Sys. Principles (SOSP)*, pages 401–416, Cascais, Portugal, Oct. 2011. Assoc. for Computing Machinery. doi: <http://doi.acm.org/10.1145/2043556.2043593>.
- [28] W. Lloyd, M. J. Freedman, M. Kaminsky, and D. G. Andersen. Stronger semantics for low-latency geo-replicated storage. In *Networked Sys. Design and Implem. (NSDI)*, pages 313–328, Lombard, IL, USA, Apr. 2013.
- [29] P. Mahajan, L. Alvisi, and M. Dahlin. Consistency, availability, and convergence. Technical Report UTCS TR-11-22, Dept. of Comp. Sc., The U. of Texas at Austin, Austin, TX, USA, 2011.
- [30] P. Mahajan, S. Setty, S. Lee, A. Clement, L. Alvisi, M. Dahlin, and M. Walfish. Depot: Cloud storage with minimal trust. *Trans. on Computer Systems*, 29(4):12:1–12:38, Dec. 2011. doi: 10.1145/2063509.2063512.
- [31] K. Petersen, M. J. Spreitzer, D. B. Terry, M. M. Theimer, and A. J. Demers. Flexible update propagation for weakly consistent replication. In *Symp. on Op. Sys. Principles (SOSP)*, pages 288–301, Saint Malo, Oct. 1997. ACM SIGOPS.
- [32] Redis. Redis is an open source, BSD licensed, advanced key-value store. <http://redis.io/>, May 2014.
- [33] N. Schiper, P. Sutra, and F. Pedone. P-Store: Genuine partial replication in wide area networks. In *Symp. on Reliable Dist. Sys. (SRDS)*, pages 214–224, New Dehli, India, Oct. 2010. IEEE Comp. Society.
- [34] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. A comprehensive study of Convergent and Commutative Replicated Data Types. Number 7506, Rocquencourt, France, Jan. 2011. URL <http://hal.archives-ouvertes.fr/inria-00555588/>.
- [35] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. Conflict-free replicated data types. In X. Défago, F. Petit, and V. Villain, editors, *Int. Symp. on Stabilization, Safety, and Security of Distributed Systems (SSS)*, volume 6976 of *Lecture Notes in Comp. Sc.*, pages 386–400, Grenoble, France, Oct. 2011. Springer-Verlag. doi: 10.1007/978-3-642-24550-3\_29.
- [36] Y. Sovran, R. Power, M. K. Aguilera, and J. Li. Transactional storage for geo-replicated systems. In *Symp. on Op. Sys. Principles (SOSP)*, pages 385–400, Cascais, Portugal, Oct. 2011. Assoc. for Computing Machinery. doi: <http://doi.acm.org/10.1145/2043556.2043592>.
- [37] 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 *Int. Conf. on Para. and Dist. Info. Sys. (PDIS)*, pages 140–149, Austin, Texas, USA, Sept. 1994.





**RESEARCH CENTRE  
PARIS – ROCQUENCOURT**

Domaine de Voluceau, - Rocquencourt  
B.P. 105 - 78153 Le Chesnay Cedex

Publisher  
Inria  
Domaine de Voluceau - Rocquencourt  
BP 105 - 78153 Le Chesnay Cedex  
inria.fr

ISSN 0249-6399