Controlled Transactional Consistency for Web Caching

Ittay Eyal   Ken Birman   Robbert van Renesse

Cornell University

Abstract

In-memory read-only caches are widely used in cloud infrastructure to reduce access latency and to reduce load on backend databases. Operators view coherent caches as impractical at genuinely large scale and many client-facing caches are updated in an asynchronous manner with best-effort pipelines.

Existing incoherent cache technologies do not support transactional data access, even if the backend database supports transactions. We propose *T-Cache*, a cache that supports read-only transactions despite asynchronous and unreliable communication with the database. We also define *cache-serializability*, a variant of serializability that is suitable for incoherent caches, and prove that with unbounded resources T-Cache implements it. With limited resources, T-Cache allows the system manager to choose a trade-off between performance and consistency.

Our evaluation shows that T-Cache detects many inconsistencies with only nominal overhead. We use synthetic workloads to demonstrate the efficacy of T-Cache when data accesses are clustered and its adaptive reaction to workload changes. With workloads based on the real-world topologies, T-Cache detects $43 - 70\%$ of the inconsistencies and increases the rate of consistent transactions by $33 - 58\%$.

1 Introduction

Internet services like online retailers and social networks store important data sets in large distributed databases. Until recently, technical challenges have forced such large-system operators to forgo transactional consistency, providing per-object consistency instead, often with some form of eventual consistency. However, recent backend systems [9, 4, 11, 10] support transactions with guarantees such as snapshot isolation and even full transactional atomicity.

Our work begins with the observation that, at present, it can be difficult for client-tier applications to leverage the transactions that the databases provide: their reads are satisfied primarily from incoherent cache. The benefits of caching are twofold. First, it reduces database load, thereby enabling higher throughput. Second, the caches are typically placed close to the clients, permitting low latency.

The problem centers on the asynchronous style of communication used between the database and the geo-distributed caches. Any approach requiring a high rate of round-trips to an authoritative backend database would cause unacceptable latency. A cache must respond instantly, and asynchronous updates rule out cache coherency schemes that would require the backend database to promptly invalidate or update cached objects, or even to track the locations at which cached objects reside. We define a variant of serializability called cache-serializability that is suitable for incoherent caches.

We then present T-Cache, a novel caching scheme that gives the application more control over consistency without changing the asynchronous coupling between caches and backend databases. We do this by storing dependency information with the cached objects, allowing the cache (or the application) to identify possible inconsistencies without contacting the database. The user can achieve a desired level of consistency by adjusting the size of this dependency data: more dependency data leads to increased consistency.

Our solution works particularly well for workloads where data accesses are clustered, which is common in today's large-scale systems. To demonstrate the efficacy of the proposed scheme, we created a prototype implementation and exposed it to workloads based on graphically-structured real-world data, such as those seen in social-networking situations. The method detects $43 - 70\%$ of the inconsistencies and can increase the ratio of consistent transactions by $33 - 58\%$, both with low overhead. We construct synthetic workloads and observe how T-Cache reacts to different clustering levels and how it adapts as clusters change.

To explain T-Cache's perfect behavior with perfectly clustered workloads we prove a related claim — we show that with unbounded resources T-Cache implements cache-serializability.

In summary, the contributions of this work are:

1. Definition of cache-serializability, a variant of serializability suitable for incoherent caches.

2. The T-Cache architecture, which allows trading off efficiency and transaction-consistency in large scale cache deployments.

3. Evaluation of T-Cache with synthetic workloads, demonstrating its adaptivity and sensitivity to clustering.

4. Evaluation of T-Cache with workloads based on graphically-structured real-world data demonstrating detection rates of $43-70\%$ and consistency improvements of $33-58\%$ with low overhead.

5. Proof that T-Cache with unbounded resources implements cache-serializability.

2 Motivation

Two-tier structure Large Internet services store vast amounts of data. Online retailers such as Amazon and eBay maintain product stocks and information, and social networking sites such as Facebook and Twitter maintain graphical databases representing user relations and group structures. For throughput, durability, and availability, such databases are sharded and replicated.

The vast majority of accesses are read-only (e.g., Facebook reports a 99.8% read rate [7]). To reduce database load and to reduce access latency, these companies employ a two-tier structure, placing layers of cache servers in front of the database (see Figure 1).

The caches of primary interest to us are typically situated far from the backend database systems — to reduce latency, companies place caches close to clients. Timeouts are used to ensure that stale cached objects will eventually be flushed, but to achieve a high cache hit ratio, timeout values are generally large. To obtain reasonable consistency, the database sends an asynchronous stream of invalidation records or cache updates, often using protocols optimized for throughput and freshness and lacking absolute guarantees of order or reliability.

Indeed, it is difficult to make this invalidation mechanism reliable without hampering database efficiency. The issues are many. First, the databases are large, residing on many servers, possibly geo-replicated. Databases use locks prudently in order to maximize concurrency. To the extent that the database keeps track of the caches that hold a copy of each object, it may be possible to send


Figure 1: The common two-tiered structure. Clients perform read-only transactions by accessing caches, which receive their values by reading from the database (solid lines). Update transactions go directly to the database (double lines). Subsequent cache invalidations can be delayed or even lost due to race conditions.

an invalidation, but tracking the state of caches is complicated and hence if they are used at all, such systems view invalidations as a kind of hint. They could be delayed (e.g., due to buffering or retransmissions after message loss), not sent (e.g., due to an inaccurate list of locations), or even lost (e.g., due to a system configuration change, buffer saturation, or because of races between reads, updates, and invalidations). A missing invalidation obviously leaves the corresponding cache entry stale. Pitfalls of such invalidation schemes are described in detail by Nishita et al. [22] and by Bronson et al. [7].

DB Transactional consistency The complexity of implementing geo-scale databases with strong guarantees initially led companies to abandon cross-object consistency altogether and make do with weak guarantees such as per-object atomicity or eventual consistency. In effect, such systems do repair any problems that arise, eventually, but the end-user is sometimes exposed to inconsistency. For some applications this is acceptable, and the approach has been surprisingly successful. In today's cloud, relaxed consistency is something of a credo.

But forgoing transactional consistency can result in undesired behavior of a service. Consider a buyer at an online site who looks for a toy train with its matching tracks just as the vendor is adding them to the database. The client may see only the train in stock but not the tracks because the product insertion transaction would often be broken into two or more atomic but independent subtransactions. In a social network, an inconsistency with unexpected results can occur if a user x's record says it belongs to a certain group, but that group's record does

not include *x*. Web albums maintain picture data and access control lists (ACLs) and it is important that ACL and album updates are consistent (the classical example involves removing one's boss from the album ACL and then adding unflattering pictures).

While many of these systems make do with weak consistency, there are certainly problems. But beyond this is a whole class of applications (such as medical ones, self-driving vehicles, management of smart buildings or the smart power grid, etc.) for which stronger consistency would appear to be an important requirement. To support such applications, there has been a wave of recent innovations within the backend, offering scalable object stores that can efficiently support transactions through snapshot isolation and even full atomicity [9, 4, 11, 10]. Our challenge is to extend the guarantees of these systems to the end-user rather than confining them primarily to the back-end database and nearby applications.

Caches cause problems As noted, today's consistency solutions are limited to the database backend. Even when the database itself is consistent, the vast majority of operations are read-only transactions issued by edge clients and are at high risk of observing inconsistent state in the cache. The outright loss of cache invalidations emerges as an especially significant problem if transactional consistency is required. An acceptable solution for a consistent cache must maintain the performance properties of the existing caching tier. First, we need to maintain the shielding role of the cache: the cache hit ratio should be high. Second, a read-only cache access should complete with a single client-to-cache round-trip on cache hits. This prohibits coherent cache solutions such as [26].

3 Architecture

Since the cache is required to respond immediately to the client on hits, and the database-cache channel is asynchronous, we decided to employ a transactional consistency that is weaker than the full ACID model. In our approach, read-only transactions and update transactions that access the same cache are guaranteed an atomic execution, but read-only transactions that access different caches may observe different orderings for independent update transactions.

Definition 1 (Cache serializability). *For every execution σ, every partial execution that includes all update transactions in σ and all read-only transactions that go through a single cache server, is serializable.*

Since it is impossible to achieve cache serializability with bounded caches and asynchronous communication with the DB, our solution seeks to approximate

it. Our idea starts with an observation: in many scenarios, objects form *clusters* with strong locality properties. Transactions are likely to access objects that are, in some sense, close to each other. For retailers this might involve related products, for social networks the set of friends, for geographical services physical proximity, and for web albums the ACL objects and the pictures assigned to them. Moreover, in some cases applications explicitly cluster their data accesses to benefit from improved parallelism [31]. The resulting transactions access objects from a single cluster, although there will also be some frequency of transactions that access unrelated objects in different clusters.

Our solution requires minor changes to the database object representation format, imposing a small and constant memory overhead (that is, independent of the database size and the transaction rate). This overhead involves tracking and caching what we refer to as *dependency lists.* These are bounded-length lists of object identifiers and the associated version numbers, each representing some recently updated objects upon which the cached object depends.

A bounded-sized list can omit dependency information required to detect inconsistencies, hence it is important to use a bound large enough to capture most of the relevant dependencies. At present we lack an automated way to do this: we require the developer to tune the length so that the frequency of errors is reduced to an acceptable level, reasoning about the trade-off (size versus accuracy) in a manner we discuss further below. Intuitively, dependency lists should be roughly the same size as the size of the workload's clusters.

Our extensions offer a transactional interface to the cache in addition to the standard read/write API. In many cases, our algorithm detects and fixes inconsistent read-only transactions at the cache with constant complexity. It does so by either aborting the transaction (which can then be retried), or invalidating a cached object which can then force a read from the database (similar to handling cache misses). When the dependency lists fail to document a necessary dependency, an application might be exposed to stale values.

Because we have in mind client-side applicatons that are unlikely to validate against the back-end, for many of our intended uses some level of undetected inconsistency can slip past. However, because the developer would often be able to tune the mechanism quite accurately, during steady-state operation of large applications, the rate of unnoticed inconsistencies could be extremely low.

With clustered workloads we will demonstrate that it is sufficient to store a small set of dependencies to detect most inconsistencies. We also investigate workloads where the clustered access pattern is less strongly evident; here, our approach is less effective even with

longer dependency list lengths. Thus our solution is not a panacea, but, for applications matched to our assumptions, can be highly effective.

3.1 Database

We assume that the database tags each object with a version number specific to the transaction that most recently updated it, and that there is a total ordering on version numbers. The version of a transaction is chosen to be larger than the versions of all objects accessed by the transaction. The database stores for each object o a list of k dependencies $(d_1^o, v_1^o), (d_2^o, v_2^o), \ldots (d_k^o, v_k^o)$. This is a list of identifiers and versions of other objects that the current version of o depends on. A read-only transaction that sees the current version of o must not see object d_i with version smaller than v_i.

When a transaction t with version v_t touches objects o_1 and o_2, it updates both their versions and their dependency lists. Subsequent accesses to object o_1 must see object o_2 with a version not smaller than v_t. Moreover, it inherits all of the l dependencies of o_2 (where l is the length of o_2's dependency list). So the dependency list of o_1 becomes

$$(d_1^{o_1}, v_1^{o_1}), (d_2^{o_1}, v_2^{o_1}), \ldots (d_k^{o_1}, v_k^{o}),$$
$$(o_2, v_t), (d_2^{o_2}, v_2^{o_2}), (d_3^{o_2}, v_3^{o_2}), \ldots (d_l^{o_2}, v_l^{o_2}) \ .$$

When a transaction is committed, this update is done for all objects in the transaction at once. Given a read set *readSet*, and a write set *writeSet*, containing tuples comprised of the keys accessed, their versions and their dependency lists, the database aggregates them to a single full dependency list as follows

$$\textit{full-dep-list} \leftarrow \bigcup_{\substack{(key,ver,depList) \in \\ readSet \cup writeSet}} \{(key,ver)\} \cup depList \ .$$

This list is pruned to match the target size using LRU, and stored with each write-set object. A list entry can be discarded if the same entry's object appears in another entry with a larger version. Nevertheless, were their lengths not bounded, dependency lists could quickly grow to include all objects in the database.

3.2 Cache

In our scheme, the cache interacts with the database in essentially the same manner as for a consistency-unaware cache, performing single-entry reads (no locks, no transactions) and receiving invalidations as the database updates objects. Unlike consistency-unaware caches, the caches read from the database not only the object's value, but also its version and the dependency list.

To its clients, the extended cache exports a transactional read-only interface. Client read requests are extended with a transaction identifier and a last-op flag

```
read(txnID, key, lastOp) .
```

The transaction identifier `txnID` allows the cache to recognize reads belonging to the same transaction. The cache responds with either the value of the requested object, or with an abort if it detects an inconsistency between this read and any of the previous reads with the same transaction ID. We do not guarantee that inconsistencies will be detected. The `lastOp` allows the cache to garbage-collect its transaction record after responding to the last read operation of the transaction. The cache will treat subsequent accesses with the same transaction ID as new transactions.

To implement this interface, the cache maintains a record of each transaction with its read values, their versions, and their dependency lists. On a read of key_{curr}, the cache first obtains the requested entry from memory (cache hit), or database (cache miss). The entry includes the value, version ver_{curr} and dependency list $depList_{curr}$. The cache checks the currently read object against each of the previously read objects. If a previously read version is older than expected by the current read's dependencies

$$\exists k, v, v' : v > v' \land (k, v) \in depList_{curr} \land$$
$$(k, v') \in readSet \cup writeSet \ , \quad (1)$$

or the current read is older than expected by the dependencies of a previous read

$$\exists v : v > v_{curr} \land (key_{curr}, v) \in readSet \cup writeSet \ , \quad (2)$$

an inconsistency is detected. Otherwise the cache returns the read value to the client.

Upon detecting an inconsistency, the cache can take one of three paths:

1. `ABORT`: abort the current transaction. Compared to the other approaches, this has the benefit of affecting only the running transaction and limiting collateral damage.

2. `EVICT`: abort the current transaction *and* evict the violating (too-old) object from the cache. This approach guesses that future transactions are likely to abort because of this object.

3. `RETRY`: check which is the violating object. If it is the currently accessed object (Equation 2), treat this access as a miss and respond to it with a value read from the database. If the violating object was returned to the user as the result of a read earlier in the transaction (Equeation 1), evict the stale object and abort the transaction.

Figure 2: Experimental setup. Update clients access database, which sends invalidations to the cache. Read-only clients access cache. Consistency monitor (experiment-only element) receives all transactions and rigorously detects inconsistencies for statistics.

4 Experimental Setup

To evaluate the effectiveness of our scheme, we implemented a prototype. To study the properties of the cache, we only need a single column (shard) of the system, namely a single cache backed by a single database server.

Figure 2 illustrates the structure of our experimental setup. A single database implements a transactional key-value store with two-phase commit. A set of cache clients perform read-only transactions through a single cache server. The cache serves the requests from its local storage if possible, or reads from the database otherwise.

On startup, the cache registers an upcall that can be used by the database to report invalidations; after each update transaction the database asynchronously sends invalidations to the cache for all objects that were modified. A ratio of 20% of the invalidations, chosen uniformly at random, are dropped by the experiment; this is extreme and would only be seen in the real world under conditions of overload or when the system configuration is changed.

Both the database and the cache report all completed transactions to a consistency monitor, created in order to gather statistics for our evaluation. This server collects both committed and aborted transactions and it maintains the full dependency graph. It performs full serialization graph testing [5] and calculates the rate of inconsistent transactions that committed and the rate of consistent transactions that were unnecessarily aborted.

Our prototype does not address the issue of cache eviction when running out of memory. In our experiments, all objects in the workload fit in the cache, and eviction is only done if there is a direct reason, as explained below. Had we modeled them, evictions would reduce the cache hit rate, but could not cause new inconsistencies.

We evaluate the effectiveness of our transactional cache using various workloads and varying the size of the dependency lists maintained by the cache and the database. For the cases considered, short dependency lists suffice (up to 5 versions). An open question for further study is whether there are workloads that might require limited but larger values.

As a baseline for comparison, we also implemented a timeout-based approach: it reduces the probability of inconsistency by limiting the life span of cache entries. We compare this method against our transactional cache by measuring its effectiveness with a varying time-to-live (TTL) for cache entries.

In all runs, both read and update transactions access 5 objects per transactions. Update clients access the database at a rate of 100 transactions per second, and read-only clients access the cache at a rate of 500 transactions per second.

Our experiment satisfies all read-only transactions from the cache, while passing all update transactions directly to the backend database. Each cache server is unaware of the other servers — it has its own clients and communicates directly with the backend database. The percentage of read-only transactions can be arbitrarily high or low in this situation: with more caches, we can push the percentage up. Our simulation focuses on just a single cache—it would behave the same had there been many cache servers.

5 Evaluation

T-Cache can be used with any transactional backend and any transactional workload. Performance for read-only transactions will be similar to non-transactional cache access: the underlying database is only accessed on cache misses. However, inconsistencies may be observed.

First, we will use synthetic workloads so we can evaluate how much inconsistency can be observed as a function of the amount of clustering in the workload. This also allows us to look at the dynamic behavior of the system, when the amount of clustering and the clustering formation change over time.

Next, we will look at workloads based on Amazon's product co-purchasing and Orkut's social network to see how much inconsistency T-Cache can detect as a function of dependency list length, and compare this with a TTL-based approach. We are also interested in overhead, particularly the additional load on the backend database that could form if the the rate of cache misses increases.

Section 3.2 presented three strategies for responding to inconsistency detection. For both the synthetic and

realistic workloads, we compare the efficacy of the three strategies.

5.1 Synthetic Workloads

Synthetic workloads allow us to understand the efficacy of T-Cache as a function of clustering. For the experiments described here, we configured T-Cache with a maximum of 5 elements per dependency list.

It might be possible for unnecessary aborts to occur. Dependency lists might come to contain irrelevant dependencies that could lead to aborts although there is no dependency. However, throughout the experiments, the rate of unnecessary aborts has been negligible.

Section 5.1.1 describes synthetic workload generation. Section 5.1.2 measures how many inconsistencies we can detect as a function of clustering and Section 5.1.3 considers clustering changes over time. Section 5.1.4 compares the efficacy of various approaches to dealing with detected inconsistencies.

5.1.1 Synthetic Workload Generation

Our basic synthetic workload is constructed as follows. We use 2000 objects numbered 0 through 1999. The objects are divided into clusters of size 5: $0-4, 5-9, 10-14, \ldots$, and there are two types of workloads. In the first, clustering is perfect and each transaction chooses a single cluster and chooses 5 times with repetitions within this cluster to establish its access set. In the second type of workloads access is not fully contained within each cluster. When a transaction starts, it chooses a cluster uniformly at random, and then picks 5 objects as follows. Each object is chosen using a bounded Pareto distribution starting at the head of its cluster i (a product of 5). If the pareto variable plus the offset results in a number outside the range (i.e., larger than 1999), the count wraps back to 0 through $i-1$.

5.1.2 Inconsistency Detection as a Function of α

We start by exploring the importance of the cluster structure by varying the α parameter of the Pareto distribution. We vary the Pareto α parameter from 1/32 to 4. In this experiment we are only interested in detection, so we choose the ABORT strategy.

Figure 3 shows the ratio of inconsistencies detected by T-Cache compared to the total number of potential inconsistencies. At $\alpha = 1/32$, the distribution is almost uniform across the object set, and the inconsistency detection ratio is low — the dependency lists are too small to hold all relevant information. At the other extreme, when $\alpha = 4$, the distribution is so spiked that almost all accesses of a transaction are within a cluster, allowing for


Figure 3: Ratio of inconsistencies as a function of α.

perfect inconsistency detection. We note that the rate of detected inconsistencies is so high at this point that much of the load goes to the backend database and saturates it, reducing the overall throughput.

5.1.3 Convergence

So far we have considered behavior with static clusters, that is, over the entire run of each experiment accesses are confined to the same (approximate) clusters. Arguably, in a real system, clusters change slowly, and so if T-Cache converges to maintain the correct dependency lists as clusters change, our setup serves as a valid quasi-static analysis.

In this section, we investigate the convergence of T-Cache when clusters change over time. The dependency lists of the objects are updated using LRU, so the dependency list of an object o tends to include those objects that are frequently accessed together with o. Dependencies in a new cluster automatically push out dependencies that are now outside the cluster.

Cluster formation

To observe convergence, we perform an experiment where accesses suddenly become clustered. Initially accesses are uniformly at random from the entire set (i.e., no clustering whatsoever), then at a single moment they become perfectly clustered into clusters of size 5. Transactions are aborted on detecting an inconsistency. We use a transaction rate of approximately 500 per second. The database includes 1000 objects.

Figure 4 shows the percentage of transactions that commit and are consistent (at the bottom), the percentage of transactions that commit but are inconsistent (in the middle), and the percentage of transactions that abort (at the top). Before $t = 58s$ access is unclustered, and as a result the dependency lists are useless; only few inconsistencies are detected, that is, about 26% of the transactions that commit have witnessed inconsistent data. At

Figure 4: Convergence of T-Cache. Before time $t = 58s$ accesses are uniformly at random. Afterward, accesses are clustered.


Figure 5: Perfectly clustered synthetic workload where the clusters shift by 1 every 3 minutes, marked by vertical lines.

$t = 58s$, accesses become perfectly clustered. As desired, we see fast improvement of inconsistency detection. The inconsistency rate drops as the abort rate rises — this is desired as well. The overall rate of consistent committed transactions drops because the probability of conflicts in the clustered scenario is higher.

Drifting Clusters

To illustrate more realistic behavior, we use clustered accesses that slowly drift. Transactions are perfectly clustered, as in the previous experiment, but every 3 minutes the cluster structure shifts by 1 ($0 - 4, 5 - 9, 10 - 14 \rightarrow 1 - 4, 5 - 10, 11 - 15$, and wrapping back to zero after 1999). Figure 5 shows the results. After each shift, the objects' dependency lists are outdated. This leads to a sudden increased inconsistency rate that converges back to zero, until this convergence is interrupted by the next shift.
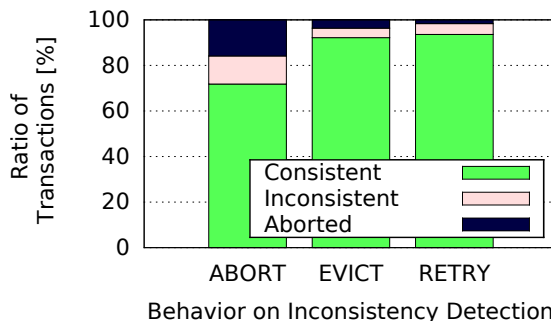

Figure 6: The efficacy of T-Cache as a function of the strategy taken for handling detected inconsistencies. ABORT detects 55% of the uncommitable tranasctions, and EVICT and RETRY reduce the rate of uncommitable transactions to about 25%.

5.1.4 Detection vs. Prevention

Section 3.2 presented three possible strategies for the cache to deal with inconsistency detection: (1) aborting the transaction (ABORT), (2) aborting and evicting value (EVICT), and (3) read-through when possible as in cache miss, abort otherwise (RETRY). We will now compare their efficacies.

We use the approximate clusters workload with 2000 objects, a window size of 5, a Pareto α parameter of 1.0, and the maximum dependency list size is set to 5.

Figure 6 illustrates the results. For each strategy, the lower portion of the graph is the ratio of committed transactions that are consistent, the middle portion is committed transactions that are inconsistent, and the top portion is aborted transactions.

The abort strategy provides a significant improvement over a normal, consistency-unaware cache, as the strategy detects and aborts over 55% of all inconsistent transactions that would have been committed. But the other strategies make further improvements. EVICT reduces uncommittable transactions to 28% of its value with ABORT. This indicates that violating (too-old) cache entries are likely to be repeat offenders: they are too old for objects that are likely to be accessed together with them in future transactions, and so it is better to evict them. RETRY reduces uncommittable transactions further to about 23% of its value with ABORT.

5.2 Realistic Workloads

We now evaluate the efficacy of T-Cache with workloads based on two sampled topologies from the online retailer Amazon and the social network Orkut. Section 5.2.1 describes how we generated these workloads. Section 5.2.2 measures the efficacy of T-Cache on these workloads as

a function of maximum dependency list size, and compares this to a strategy based on TTLs. Section 5.2.3 compares the efficacy of the three strategies of dealing with detected inconsistencies.

5.2.1 Workload Generation

We generated two workloads based on real data:

1. Amazon: We started from a snapshot of Amazon's product co-purchasing graph taken early 2003 [15]. Each product sold by the online retailer is a node and each pair of products that were purchased in a single user session is an edge. The original graph contains more than 260,000 nodes.

2. Orkut: For the second, we used a snapshot of the friendship relations graph in the Orkut social network, taken late 2006 [21]. In this graph, each user is a node and each pair of users that have a friend relationship is an edge. The original graph contains more than 3,000,000 nodes.

Because the sampled topologies are large and we only need to simulate a single "column" of the system for our purposes — one database server and one cache server — we down-sample both graphs to 1000 nodes. We use a technique based on random walks that maintains important properties of the original graph [16], specifically clustering which is central to our experiment. We start by choosing a node uniformly and random and start a random walk from that location. In every step, with probability 15%, the walk reverts back to the first node and start again. This is repeated until the target number of nodes have been visited. Figure 7(a) and (b) show a further down-sampling to 500 nodes to provide some perception of the topologies. The graphs are visibly clustered, the Amazon topology more so than the Orkut one, yet well-connected.

Treating nodes of the graphs as database objects, transactions are likely to access objects that are topologically close to one another. For the online retailer, it is likely that objects bought together are also viewed and updated together (e.g., viewing and buying a toy train and matching rails). For the social network, it is likely that data of befriended users are viewed and updated together (e.g., tagging a person in a picture, commenting on a post by a friend's friend, or viewing one's neighborhood).

Therefore, we generate a transactional workload that accesses products that are topologically close. Again, we use random walks. Each transaction starts by picking a node uniformly at random and takes 5 steps of a random walk. The nodes visited by the random walk are the objects the transaction accesses. Update transactions first read all objects from the database, and then update all objects at the database. Read transactions read the objects directly from the cache.

5.2.2 Efficacy and Overhead

In this section we evaluate T-Cache using the workloads described above. We found that the abort rate is negligible in all runs. Efficacy is therefore defined to be the ratio of inconsistent transactions out of all commits.

The overhead of the system is twofold. First, dependency list maintenance implies storage and bandwidth overhead at both the database and the cache, as well as compute overhead for dependency list merging at the server and consistency checks at the cache. However, the storage required is only for object IDs and versions, not content, and both updates and checks are $O(1)$ in the number of objects in the system and $O(k^2)$ in the size of the dependency lists, which is limited to 5 in our experiments.

The second and potentially more significant overhead is the effect on cache hit ratio due to evictions and hence the database load. Since cache load is significantly larger than database load (2 orders of magnitude for Facebook [7]), even a minor deterioration in hit ratio can yield a prohibitive load on the backend database. Figure 7c shows the experiment results. Each data point is the result of a single run.

We vary the dependency list size and for each value run the experiment for the two workloads and measure the average values of these metrics. T-Cache is able to reduce inconsistencies significantly. For the retailer workload, a single dependency reduces inconsistencies to 56% of their original value, two dependencies reduce inconsistencies to 11% of their original value, and three to less than 7%. For the social network workload, with 3 dependencies fewer than 7% of the inconsistencies remain.

In both workloads there is no visible effect on cache hit ratio, and hence no increased access rate at the database. The reduction in inconsistency ratio is significantly better for the retailer workload. Its topology has a more clustered structure, and so the dependency lists hold more relevant information.

Next we compared our technique with a simple approach in which we limited the life span (Time To Live, TTL) of cache entries. Here inconsistencies are not detected, but their probability of being witnessed is reduced by having the cache evict entries after a certain period even if the database did not indicate they are invalid.

We run a set of experiments similar to the T-Cache ones, varying cache entry TTL to evaluate the efficacy of this method in reducing inconsistencies and the corresponding overhead. Compared to T-Cache, Limiting TTL has detrimental effects on cache hit ratio,

(a) Product Affinity (Amazon)

(b) Social Network (Orkut)

(c) Transactional Cache

(d) Limited Cache Entry TTL

Figure 7: Experiments with workloads based on a web retailer product affinity topology and a social network topology illustrated in (a) and (b). Transactional cache (c) compared against the alternative of reducing cache entry time-to-live (d). Data points are medians and error bars bound the 10 and 90 percentiles.
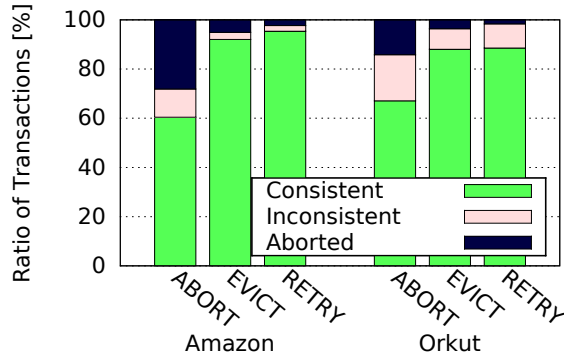
Figure 8: The efficacy of T-Cache as a function of the inconsistency handling strategy for realistic workloads.

quickly increasing the database workload. By increasing database access rate to more than twice its original load we only observe a reduction of inconsistencies of about 10%. This is more than twice the rate of inconsistencies achieved by T-Cache for the retailer workload and only slightly better than the rate of inconsistencies achieved by T-Cache for the social network workload; and with twice the additional load on the database.

5.2.3 Detection vs. Prevention

Figure 8 compares the efficacy of the ABORT, EVICT and RETRY policies with the Amazon and Orkut workloads. In these experiments we use dependency lists of length 3. Just as with the synthetic workload, evicting conflicting transactions is an effective way of invalidating stale objects that might cause problems for future transactions.

The effects are more pronounced for the well-clustered Amazon workload. With the Amazon workload, ABORT is able to detect 70% of the inconsistent transactions, whereas with the less-clustered Orkut workload it only detects 43%. In both cases ABORT reduces uncommittable transactions considerably, relative to their value with ABORT — 20% with the Amazon workload and 36% with Orkut. In the Amazon workload, RETRY further reduces this value to 11% of its value with ABORT.

6 Consistency

As demonstrated in Section 5.1.3, T-Cache appears to converge to perfect detection when stable clusters that are as large as its dependency lists. In such a scenario, the dependency lists are large enough to describe all relevant dependencies. To show that this indeed leads to perfect detection we prove the following theorem.

Theorem 1. *T-Cache with unbounded cache size and unbounded dependency lists implements cache-serializability.*

Since we assume that the transactional DB is serializable, the operations in an execution of update transactions σ_{update} can be serialized as some serial execution π. The next claim trivially follows from the definition of the database dependency list specification:

Claim 1. *If π is a serialization of the update transactions of an execution σ_{update}, then, at every step in π, the version dependencies of every object match those stored in its dependency list.*

To prove Theorem 1, we first describe a routine for placing a read-only transaction from a cache server in a serialization of a subset of σ, to form a serialization of both the update transaction and the read-only transaction.

6.1 Permutation routine

Let σ be an execution of the T-Cache system, and denote by σ_{update} the projection of σ on the set of database update transactions. Transaction T reads objects $o_1, o_2, \ldots, o_n$ with versions $v_1, v_2, \ldots, v_n$, respectively. Take any serialization π of σ_{update} (one exists according to Claim 1) and consider the first time when all the objects the transaction reads are at a version at least as large as the versions that T reads. At this time at least one object read by T, the last written according to π, has the correct version, but others might not. Assume WLOG that the last version written is v_n of object o_n at step t of π. Denote by t' the latest time at which a wrong version (not the one read by T) is written, and assume WLOG it is version $v_{n-1} + k$ of object o_{n-1} (rather than the desired version v_{n-1}) for some $k \geq 1$.

We now describe a single step of the routine. Consider the transactions between t' and t (inclusive). Divide these transactions into three sets:

Set 1 Transactions dependent on the transaction at t' (including t').

Set 2 Transactions on which t is dependent (including t).

Set 3 Transactions that do not belong to either group.

The following Lemma states that there is no dependency among objects in sets 1 and 2, and hence there is no intersection between the sets.

Lemma 1. *Sets 1 and 2 are independent.*

Proof. If they were dependent, then version v_n of object o_n depends on version $v_{n-1} + k$ of object o_{n-1}, and this dependency is reflected in their T-Cache dependency

lists, because they are unbounded. However, transaction T has read version v_{n-1} of object o_{n-1}, which is older than $v_{n-1} + k$. The read of the stale version v_{n-1} of o_{n-1} would have been detected by T-Cache and the transaction would have been aborted. Therefore the assumption is wrong, and the sets are indeed independent. $\square$

Set 3, perhaps an empty set, is unrelated to sets 1 and 2 by definition. We therefore switch sets 1 and 2, and place set 3 right after them, maintaining a serialization of σ_{update}.

For example, consider the following serialization: (X_i denotes a transaction X in set i):

A_1	B_3	C_1	D_1	E_3	F_2	G_2

After the permutation, we obtain:

A_1	C_1	D_1	F_2	G_2	B_3	E_3

Performing this permutation is one step of the routine. We repeat this step forming a series of permutations. Each permutation is a serialization of σ_{update}, and each permutes a range of the transactions with respect to the previous step. In each step the right end of the range is earlier than in the previous step, as one or more of the objects is closer to the value read by T. Eventually we therefore reach a permutation where at the chosen time all read objects are at their correct versions. We place T there to obtain the desired serialization of the update transactions and T.

6.2 T-Cache Consistency

We proceed to prove Theorem 1.

Proof. Let σ be an execution of the T-Cache system, and denote by σ_{update} the projection of σ on the set of database update transactions. Denote by $T_1, T_2, \ldots, T_m$ a set of read-only transactions performed through a single T-Cache server.

If the read sets of two transactions include the same object o, we say the one that read a larger version of o depends on the other. All transactions access the same cache, and the cache is unbounded. Therefore, values are only replaced by newer versions, so it is easy to see that there are no cycles such that two transactions depend on one another. The dependency graph therefore describes a partial order of the read-only transactions, and we choose an arbitrary total ordering that respects this partial order. Assume WLOG the order is $T_1, T_2, \ldots, T_m$.

We take an initial arbitrary serialization π_0 of σ and permute it according to the route above to place T_1, the first read-only transaction. The result is a permutation π'_1 that includes T_1. Then, we take all transactions that precede T_1 in π'_1 although T_1 does not depend on them, and place them after T_1. We call this permutation π_1.

Next we place T_2 by permuting π_1. If T_2 can be placed immediately after T_1, we place it there to form π_2. If T_2 is independent of T_1 then all its preceding transactions (according to the dependency graph) are unrelated to T_1 and are therefore located after it. The permutations required are therefore after T_1's location. Finally, if T_2 depends on T_1, all relevant update transactions are located after T_1 in π_1, and therefore the permutations required are all after T_1's location. Since in all cases the permutations are after T_1's location in π_1, they do not affect the correctness of T_1's placement. We take the resulting permutation that we call π'_2, and move all transactions that neither T_2 nor T_1 depend on to right after T_2. The resulting permutation is π_2.

We repeat this process until we place all read-only transactions, forming π_m. This is a serialization of the update transactions in σ and all read-only transactions that accessed the same cache. We have therefore shown that in any execution of T-Cache the update transactions can be serialized with read-only transactions that accessed a single cache, which means that T-Cache implements cache serializability. $\square$

7 Related Work

Scalable Consistent Databases Recent years have seen a surge of progress in the development of scalable object stores that support transactions. Some systems such as [18, 27, 17, 30] export novel consistency definitions that allow for effective optimizations. Several recent systems implement full fledged atomicity while preserving the system's scalability with a wide variety of workloads. Google's Spanner utilizes accurate clock synchronization. Tango [4] by Balakrishnan et al. is constructed on top of the scalable Corfu [3] log. Eyal et al. [11] utilize a large set of independent logs. Escriva et al. [10] use DHT-based locking. Zhang et al. [33] use lock chains and assume transactions are known in advance. These methods all scale well and in many cases allow databases to accept loads similar to those handled by non-transactional databases. Nevertheless, they are not expected to disrupt the prevailing two-tier structure; *caches remain invaluable*.

Consistent Caching Much work has been done on creating consistent caches for web servers [8, 32, 34, 1, 23], distributed file systems [13, 29], Key-Value Stores [20, 7, 22] and higher level objects [12, 2]. Such systems consider only one object at a time, and only individual read and write operations, as they do not support a transactional interface. There are few if any multi-object or multi-operation consistency considerations. These systems generally try to avoid staleness through techniques

such as Time-To-Live (TTL), invalidation broadcasts, and leases. *Our work considers multi-object transactional consistency of cache access*.

Transactional Caching Early work on scalable database caching mostly ignored transactional consistency [19]. Since then, work has been done on creating consistent caches for databases. TxCache [26] extends a centralized database with support for caches that provide snapshot isolation semantics, albeit the snapshots seen may be stale. To improve the commit rate for read-only transactions, they use *multiversioning*, where the cache holds several versions of an object and enables the cache to choose a version that allows a transaction to commit. This technique could also be used with our solution. Perez-Sorrosal et al. [24, 25] also support snapshot isolation, but can be used with any backend database, including ones that are sharded and/or replicated. JBossCache [28] provides a transactionally consistent cache for the JBoss middleware. Both JBossCache and [14] support transactions on cached Enterprise JavaBeans. [6] allows update transactions to read stale data out of caches and provide bounds on how much staleness is allowed. These techniques require fast communication between the cache and the database for good performance. In contrast, *in our work caches are asynchronously updated (or invalidated)*, which is how caches currently work in large multi-regional clouds.

8 Future Directions

The dependency list sizes for all objects in T-Cache are currently all of the same maximum length. This may not be optimal. For example, if the workload accesses objects in clusters of different sizes, objects of larger clusters call for longer dependency lists. Once appropriate real workloads are available, it may be possible to improve performance by dynamically changing per-object dependency list sizes, balancing between objects to maintain the same overall space overhead. Another option is to explore an approach in which each type of object would have its own dependency list bound; this could work well if a system has multiple classes of objects, all clustered but with different associated clustering properties.

At present, T-Cache is semantics-agnostic and treats all objects and object relations as equal, using an LRU policy to trim the list of dependencies. However, there may be cases in which the application could explicitly inform the cache of relevant object dependencies, and those could then be treated as more important and retained, while other less important ones are managed by some other policy such as LRU. For example, in a web album the set of pictures and their ACL is an important dependency whereas occasional tagging operations that relate pictures to users may be less important. It may be straightforward to extend the cache API to allow the application to specify such dependencies and to modify T-Cache to respect them.

9 Conclusion

Existing large-scale computing frameworks make heavy use of edge caches to reduce client latency, but this form of caching has not been available for transactional applications. We believe this is one reason that transactions are generally not considered to be a viable option in extremely large systems.

We defined cache-serializability, a variant of serializability that is suitable for incoherent caches. We then presented T-Cache, an architecture for controlling transaction consistency with caches. The system extends the edge cache by allowing it to offer a transactional interface. We believe that T-Cache is the first transaction-aware caching architecture in which caches are updated asynchronously. In particular, a lookup request only requires a round-trip to the database in case there is a cache miss — there is no additional traffic and delays to ensure cache coherence.

T-Cache associates dependency information with cached database objects, while leaving the interaction between the backend systems and the cache otherwise unchanged. This information includes version identifiers and bounded-length dependency lists. With this modest amount of additional information, we show that inconsistency can be greatly reduced or even completely eliminated in some cases.

T-Cache is intended for clustered workloads, and those arise naturally in social networks, product relationships, mobile applications with spatial locality, and so on. Our experiments demonstrate T-Cache to be effective in realistic workloads based on datasets from Amazon and Orkut. Using dependency lists of size 3, T-Cache detected $43 - 70\%$ of the inconsistencies, and was also able to increase consistent transaction rate by $33 - 58\%$ with only nominal overhead on the database. Our experiments with synthetic workloads showed that T-Cache's efficacy depends on the clustering level of the workload. T-Cache adapts to dynamically changing workloads where clusters change over time.

Due to resource limitations T-Cache maintains only a short dependency list, which is naturally imperfect and does not include all dependencies. We proved that when resources are unbounded, T-Cache's algorithm implements cache-serializability.

References

[1] M.H.S. Attar and M.T. Ozsu. Alternative architectures and protocols for providing strong consistency in dynamic web applications. *World Wide Web Journal*, 9(3):215–251, 2006.

[2] R. Bakalova, A. Chow, C. Fricano, P. Jain, N. Kodali, D. Poirier, S. Sankaran, and D. Shupp. WebSphere dynamic cache: Improving J2EE application experience. *IBM Systems Journal*, 43(2), 2004.

[3] Mahesh Balakrishnan, Dahlia Malkhi, Vijayan Prabhakaran, Ted Wobber, Michael Wei, and John D Davis. Corfu: A shared log design for flash clusters. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI'12)*, pages 1–14, 2012.

[4] Mahesh Balakrishnan, Dahlia Malkhi, Ted Wobber, Ming Wu, Vijayan Prabhakaran, Michael Wei, John D Davis, Sriram Rao, Tao Zou, and Aviad Zuck. Tango: Distributed data structures over a shared log. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles*, pages 325–340. ACM, 2013.

[5] Philip A Bernstein. *Concurrency control and recovery in database systems*, volume 370. Addison-Wesley, New York, 1987.

[6] Philip A. Bernstein, Alan Fekete, Hongfei Guo, Raghu Ramakrishnan, and Pradeep Tamma. Relaxed-currency serializability for middle-tier caching and replication. In *International Conference on Management of Data (SIGMOD)*, pages 599–610, 2006.

[7] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, Harry C Li, et al. TAO: Facebook's distributed data store for the social graph. In *USENIX Annual Technical Conference*, pages 49–60, 2013.

[8] J. Challenger, A. Iyengar, and P. Dantzig. A scalable system for consistently caching dynamic web data. In *Proc. INFOCOM 99*, March 1999.

[9] James C Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, et al. Spanner: Googles globally distributed database. *ACM Transactions on Computer Systems (TOCS)*, 31(3):8, 2013.

[10] Robert Escriva, Bernard Wong, and Emin Gün Sirer. Warp: Multi-key transactions for key-value stores. Technical report, Dept. of Computer Science, Cornell University, 2013.

[11] Ittay Eyal, Ken Birman, Idit Keidar, and Robbert van Renesse. Ordering transactions with prediction in distributed object stores. In *Proc. of the 7th Workshop on Large-Scale Distributed Systems and Middleware (LADIS'13)*, 2013.

[12] Steven D. Gribble, Eric A. Brewer, Joseph M. Hellerstein, and David Culler. Scalable, distributed data structures for internet service construction. In *Proceedings of the 4th Conference on Symposium on Operating System Design & Implementation (OSDI'00)*, volume 4. USENIX Association, 2000.

[13] Christopher Angel Kent. *Cache Coherence in Distributed Systems*. PhD thesis, Purdue University, August 1986.

[14] A. Leff and J.T. Rayfield. Improving application throughput with enterprise JavaBeans caching. In *International Conference on Distributed Computing Systems (ICDCS)*, 2003.

[15] Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.

[16] Jure Leskovec and Christos Faloutsos. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 631–636. ACM, 2006.

[17] Cheng Li, Daniel Porto, Allen Clement, Johannes Gehrke, Nuno Preguica, and Rodrigo Rodrigues. Making geo-replicated systems fast as possible, consistent when necessary. In *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI '12)*, 2012.

[18] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual: Scalable causal consistency for wide-area storage with COPS. In *Proc. 23rd ACM Symposium on Operating Systems Principles (SOSP 11)*, October 2011.

[19] Qiong Luo, Sailesh Krishnamurthy, C. Mohan, Hamid Pirahesh, Honguk Woo, Bruce G. Lindsay, and Jeffrey F. Naughton. Middle-tier database caching for e-business. In *International Conference on Management of Data (SIGMOD)*, pages 600–611, 2002.

[20] Memcached. Memcached: a distributed memory object caching system. http://memcached.org.

[21] Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC'07)*, pages 29–42. ACM, 2007.

[22] Rajesh Nishtala, Hans Fugal, Steven Grimm, Marc Kwiatkowski, Herman Lee, Harry C. Li, Ryan McElroy, Mike Paleczny, Daniel Peek, Paul Saab, David Stafford, Tony Tung, and Venkateshwaran Venkataramani. Scaling Memcache at Facebook. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI'13)*, pages 385–398, Lombard, IL, 2013. USENIX.

[23] Oracle. Oracle web cache. http://www.oracle.com/technetwork/middleware/webtier/overview.

[24] Francisco Perez-Sorrosal, Marta Patino-Martinez, Ricardo Jimenez-Peris, and Bettina Kemme. Consistent and scalable cache replication for multi-tier J2EE applications. In *Proc. of Middleware'07*, 2007.

[25] Francisco Perez-Sorrosal, Marta Patiño-Martinez, Ricardo Jimenez-Peris, and Bettina Kemme. Elastic SI-Cache: consistent and scalable caching in multi-tier architectures. *The International Journal on Very Large Data Bases (VLDB Journal)*, 20(6):841–865, 2011.

[26] Dan R.K. Ports, Austin T. Clements, Irene Zhang, Samuel Madden, and Barbara Liskov. Transactional consistency and automatic management in an application data cache. In *9th USENIX Symposium on Operating Systems Design and Implementation (OSDI '10)*, volume 10, pages 1–15, 2010.

[27] Yair Sovran, Russell Power, Marcos K. Aguilera, and Jinyang Li. Transactional storage for geo-replicated systems. In *Proc. 23rd ACM Symposium on Operating Systems Principles (SOSP 11)*, October 2011.

[28] Manik Surtani and Bela Ban. JBoss Cache. http://jbosscache.jboss.com.

[29] Amin M. Vahdat, Paul C. Eastham, and Thomas E. Anderson. Webfs: A global cache coherent file system. Technical report, UC Berkeley, December 1996.

[30] Chao Xie, Chunzhi Su, Manos Kapritsos, Yang Wang, Navid Yaghmazadeh, Lorenzo Alvisi, and Prince Mahajan. Salt: Combining ACID and BASE in a distributed database. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI '14)*, 2014.

[31] Wenlei Xie, Guozhang Wang, David Bindel, Alan Demers, and Johannes Gehrke. Fast iterative graph computation with block updates. In *Proc. of the VLDB Endowment (PVLDB)*, volume 6, pages 2014–2025, September 2013.

[32] H. Yu, L. Breslau, and S. Shenker. A scalable web cache consistency architecture. *SIGCOMM Computer Communications Review*, 29(4):163–174, 1999.

[33] Yang Zhang, Russell Power, Siyuan Zhou, Yair Sovran, Marcos K Aguilera, and Jinyang Li. Transaction chains: achieving serializability with low latency in geo-distributed storage systems. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles*, pages 276–291. ACM, 2013.

[34] H. Zhu and T. Yang. Class-based cache management for dynamic web content. In *Proc. INFOCOM 01*, 2001.