

Morphus: Supporting Online Reconfigurations in Sharded NoSQL Systems

Mainak Ghosh, Wenting Wang, Gopalakrishna Holla, Indranil Gupta
Department of Computer Science
University of Illinois, Urbana Champaign
{mghosh4, wwang84, hollava2, indy}@illinois.edu

Abstract—While sharded NoSQL stores offer high availability, reconfiguration operations present a major pain point in deployments today. For instance, in order to change a configuration setting such as the shard (or primary) key of a database table, the prevalent solutions entail shutting down the database, exporting and re-importing the table, and restarting the database. This goes against the NoSQL philosophy of high availability of data.

Our system, called Morphus, provides support towards reconfigurations for NoSQL stores in an *online* manner. Morphus allows read and write operations to continue concurrently with the data transfer among servers. Morphus works for NoSQL stores that feature master-slave replication, range partitioning, and flexible data placement. This paper presents: i) a systems architecture for online reconfigurations, incorporated into MongoDB, and ii) optimal algorithms for online reconfigurations. Our evaluation using realistic workloads shows that Morphus completes reconfiguration efficiently, offers high availability, and incurs low overhead for reads and writes.

Index Terms—shard key, reconfiguration, key value stores

1 INTRODUCTION

Distributed NoSQL storage systems comprise one of the core technologies in today’s cloud computing revolution. These systems are attractive because they offer high availability and fast read/write operations for data. They are used in production deployments for online shopping, content management, archiving, e-commerce, education, finance, gaming, email and healthcare. The NoSQL market is expected to earn \$14 Billion revenue during 2015-2020, and become a \$3.4 Billion market by 2020 [16].

In production deployments of NoSQL systems and key-value stores, *reconfiguration operations* have been a persistent and major pain point [3], [9]. (Even traditional databases have considered reconfigurations to be a major problem over the past two decades [14], [42].) Reconfiguration operations deal with changes to configuration parameters at the level of a database table or the entire database itself, in a way that affects a large amount of data all at once. Examples include schema changes like changing the shard/primary key which is used to split a table into blocks, or changing the block size itself, or increasing or decreasing the cluster size (scale out/in).

In today’s sharded NoSQL deployments [8], [11], [24], [37], such data-centric¹ global reconfiguration operations are quite inefficient. This is because executing them relies on ad-hoc mechanisms rather than solving the core underlying algorithmic and system design problems. The most common solution involves first saving a table or the entire database, and then re-importing all the data into the new configuration [2]. This approach leads to a significant period of unavailability. A second option may be to create a new cluster of servers with the new database configuration and then populate it with data from the old cluster [21], [31], [33]. This approach does not support concurrent reads and writes during migration, a feature we would like to provide.

Consider an admin who wishes to change the shard key inside a sharded NoSQL store like MongoDB [2]. The shard key is used to split the database into blocks, where each block stores values for a contiguous range of shard keys. Queries are answered much faster if they include the shard key as a query parameter (otherwise the query needs to be multicast). Today’s systems strongly recommend that the admin decide the shard key at database creation time, but not change it afterwards. However, this is challenging because it is hard to guess how the workload will evolve in the future.

As a result, many admins start their databases with a system-generated UID as the shard key, while others hedge their bets by inserting a surrogate key with each record so that it can be later used as the new primary key. The former UID approach reduces the utility of the primary key to human users and restricts query expressiveness, while the latter approach would work only with a good guess for the surrogate key that holds over many years of operation. In either case, as the workload patterns become clearer over a longer period of operation, a new application-specific shard key (e.g., username, blog URL, etc.) may become more ideal than the originally-chosen shard key.

Broadly speaking, admins need to change the shard key prompted by either changes in the nature of the data being received, or due to evolving business logic, or by the need to perform operations like join with other tables, or due to prior design choices that are sub-optimal in hindsight. Failure to reconfigure databases can lower performance, and also lead to outages such

This work was supported in part by the following grants: NSF CNS 1319527, NSF CCF 0964471, NSF CNS 1409416 and AFOSR/AFRL FA8750-11-2-0084.

1. Data-centric reconfiguration operations deal only with migration of data residing in database tables. Non-data-centric reconfigurations are beyond our scope, e.g., software updates, configuration table changes, etc.

as the one at Foursquare [15]. As a result, the reconfiguration problem has been a fervent point of discussion in the community for many years [1], [7]. While such reconfigurations may not be very frequent operations, they are a significant enough pain point that they have engendered multiple JIRA (bug tracking) entries (e.g., [3]), discussions and blogs (e.g., [7]). Inside Google, resharding MySQL databases took two years and involved a dozen teams [29]. Thus, we believe that this is a major problem that directly affects the life of a system administrator – without an automated reconfiguration primitive, reconfiguration operations today are laborious and manual, consume significant amounts of time, and leave open the room for human errors during the reconfiguration.

In this paper, we present a system that supports automated reconfiguration. Our system, called Morpheus, allows reconfiguration changes to happen in an *online* manner, that is, by concurrently supporting reads and writes on the database table while its data is being reconfigured.

Morpheus assumes that the NoSQL system features: 1) master-slave replication, 2) range-based sharding (as opposed to hash-based)², and 3) flexibility in data assignment³. Several databases satisfy these assumptions, e.g., MongoDB [11], RethinkDB [13], CouchDB [4], etc. To integrate our Morpheus system we chose MongoDB due to its clean documentation, strong user base, and significant development activity. To simplify discussion, we assume a single datacenter, but later present geo-distributed experiments. Finally, we focus on NoSQL rather than ACID databases because the simplified CRUD (Create, Read, Update, Delete) operations allow us to focus on the reconfiguration problem – addressing ACID transactions is an exciting avenue that our paper opens up.

Morpheus solves three major challenges: 1) in order to be fast, data migration across servers must incur the least traffic volume; 2) degradation of read and write latencies during reconfiguration must be small compared to operation latencies when there is no reconfiguration; 3) data migration traffic must adapt itself to the datacenter’s network topology.

We solve these challenges via these approaches:

- Morpheus performs automated reconfiguration among cluster machines, without needing additional machines.
- Morpheus concurrently allows efficient read and write operations on the data that is being migrated.
- To minimize data migration volume, Morpheus uses new algorithms for placement of new data blocks at servers.
- Morpheus is topology-aware by using new network-aware techniques which adapt data migration flow as a function of the underlying network latencies.
- We integrate Morpheus into MongoDB [2] focusing on the shard key change operation.
- Our experiments using real datasets and workloads show that Morpheus maintains several 9s of availability during reconfigurations, keeps read and write latencies low, and scales well with both cluster and data sizes.

2. Most systems that hash keys use range-based sharding of the hashed keys, and so our system applies there as well. Our system also works with pure hash sharded systems, though it is less effective.

3. This flexibility allows us to innovate on data placement strategies. Inflexibility in consistent hashed systems like Cassandra [37] require a different solution, which is the target of a different project of ours.




Fig. 1: **Morpheus Phases. Arrows represent RPCs. M stands for Master, S for Slave.**

2 SYSTEM DESIGN

We now describe the design of our Morpheus system.

2.1 MongoDB System Model

We have chosen to incorporate Morpheus into a popular sharded key-value store, MongoDB [11] v2.4. Beside its popularity, our choice of MongoDB is also driven by its clean documentation, strong user base, and significant development and discussion around it.

A MongoDB deployment consists of three types of servers. The *mongod* servers store the data chunks themselves – typically, they are grouped into disjoint *replica sets*. Each replica set contains the same number of (typically 3) servers which are exact replicas of each other, with one of the servers marked as a primary (master), and others acting as secondaries (slaves). The configuration parameters of the database are stored at the *config servers*. Clients send CRUD (Create, Read, Update, Delete) queries to a front-end server, called *mongos*. The *mongos* servers also cache some of the configuration information from the *config servers*, e.g., in order to route queries they cache mappings from each chunk range to a replica set.

A single database table in MongoDB is called a *collection*. Thus, a single MongoDB deployment consists of several collections.

2.2 Reconfiguration Phases in Morpheus

Section 3 will describe how new chunks can be allocated to servers in an optimal way. Given such an allocation, we now describe how to i) migrate data, while ii) concurrently supporting operations on this data.

Overview. Morphus allows a reconfiguration operation to be initiated by a system administrator on any collection. Morphus executes the reconfiguration via five sequential phases, as shown in Figure 1.

First Morphus prepares for the reconfiguration by creating partitions (with empty new chunks) using the new shard key (Prepare phase). Second, Morphus isolates one secondary server from each replica set (Isolation phase). In the third Execution phase, these secondaries exchange data based on the placement plan decided by mongos. In the meantime, further operations may have arrived at the primary servers – these are now replayed at the secondaries in the fourth Recovery phase. When the reconfigured secondaries are caught up, they swap places with their primaries (Commit phase).

At this point, the database has been reconfigured and can start serving queries with the new configuration. However, other secondaries in all replica sets need to reconfigure as well – this slave catchup is done in multiple rounds, with the number of rounds equal to the size of the replica set.

Next we discuss the individual phases in detail.

Prepare. The first phase is the Prepare phase, which runs at the mongos front-end. Reads and writes are not affected in this phase, and can continue normally. Concretely, for the shard key change reconfiguration, there are two important steps in this phase:

- **Create New Chunks:** Morphus queries one of the mongod servers and sets split points for new chunks by using MongoDB’s internal splitting algorithm (for modularity).
- **Disable Background Processes:** We disable background processes of the NoSQL system which may interfere with the reconfiguration transfer. This includes the MongoDB Balancer, a background thread that periodically checks and balances the number of chunks across replica sets.

Isolation. In order to continue serving operations while the data is being reconfigured, Morphus first performs reconfiguration transfers only among secondary servers, one from each replica set. It prepares by performing two steps:

- **Mute Slave Olog Replay:** Normally, the primary server forwards the operation log (called *oplog*) of all the write operations it receives to the secondary, which then replays it. In the Isolation phase, this oplog replay is disabled at the selected secondaries, but only for the collection being reconfigured – other collections still perform oplog replay. We chose to keep the secondaries isolated, rather than removed, because the latter would make Recovery more challenging by involving collections not being resharded.
- **Collect Timestamp:** In order to know where to restart replaying the oplog in the future, the latest timestamp from the oplog is stored in memory by mongos.

Execution. This phase is responsible for making placement decisions (which we will describe in Section 3) and executing the resultant data transfer among the secondaries. In the meantime, the primary servers concurrently continue to serve client CRUD operations. Since the selected secondaries are isolated, a consistent snapshot of the collection can be used to run the placement algorithms (Section 3) at a mongos server.

Assigning a new chunk to a mongod server implies migrating data in that chunk range from several other mongod servers to the assigned server. For each new chunk, the assigned server creates a separate TCP connection to each source server, “pulls” data

from the appropriate old chunks, and commits it locally. All these migrations occur in parallel. We call this scheme of assigning a single socket to each migration as “chunk-based”. The chunk-based strategy can create stragglers, and Section 4.1 addresses this problem.

Recovery. At the end of the Execution phase, the secondary servers have data stored according to the new configuration. However, any write (create, update or delete) operations that had been received by a primary server, since the time its secondary was isolated, now need to be communicated to the secondary.

Each primary forwards each item in the oplog to its appropriate new secondary, based on the new chunk ranges. This secondary can be located from our placement plan in the Execution phase, if the operation involved the new shard key. If the operation does not involve the new shard key, it is multicast to all secondaries, and each in turn checks whether it needs to apply it. This mirrors the way MongoDB typically routes queries among replica sets.

However, oplog replay is an iterative process – during the above oplog replay, further write operations may arrive at the primary oplogs. Thus, in the next iteration this delta oplog will need to be replayed. If the collection is hot, then these iterations may take very long to converge. To ensure convergence, we adopt two techniques: i) cap the replay at 2 iterations, and ii) enforce a write throttle before the last iteration. The write throttle is required because of the atomic commit phase that follows right afterwards. The write throttle rejects any writes received during the final iteration of oplog replay. An alternative was to buffer these writes temporarily at the primary and apply them later – however, this would have created another oplog and reinstated the original problem. In any case, the next phase (Commit) requires a write throttle anyway, and thus our approach dovetails with the Commit phase. Read operations remain unaffected and continue normally.

Commit. Finally, we bring the new configuration to the forefront and install it as the default configuration for the collection. This is done in one atomic step by continuing the write throttle from the Recovery phase.

This atomic step consists of two substeps:

- **Update Config:** The mongos server updates the config database with the new shard key and chunk ranges. Subsequent CRUD operations use the new configuration.
- **Elect Slave As Master:** Now the reconfigured secondary servers become the new primaries. The old primary steps down and Morphus ensures the secondary wins the subsequent leader election inside each replica set.

To end this phase, the new primaries (old secondaries, now reconfigured) unmute their oplog and the new secondaries (old primaries for each replica set, not yet reconfigured) unthrottle their writes.

Read-Write Behavior. The end of the Commit phase marks the switch to the new shard key. Until this point, all queries with old shard key were routed to the mapped server and all queries with new shard key were multicast to all the servers (normal MongoDB behavior). After the Commit phase, a query with the new shard key is routed to the appropriate server (new primary). Queries which do not use the new shard key are handled with a multicast, which is again normal MongoDB behavior.

Reads in MongoDB offer per-key sequential consistency. Morphus is designed so that it continues to offer the same consistency model for data under migration.

Slave Isolation & Catchup. After the Commit phase, the secondaries have data in the old configuration, while the primaries receive writes in the new configuration. As a result, normal oplog replay cannot be done from a primary to its secondaries. Thus, Morplus isolates all the remaining secondaries simultaneously.

The isolated secondaries catch up to the new configuration via (*replica set size* - 1) sequential rounds. Each round consists of the Execution, Recovery and Commit phases shown in Figure 1. However, some steps in these phases are skipped – these include the leader election protocol and config database update. Each replica set numbers its secondaries and in the i th round ($2 \leq i \leq \text{replica set size}$), its i th secondary participates in the reconfiguration. The group of i th secondaries reuses the old placement decisions from the first round’s Execution phase – we do so because secondaries need to mirror primaries.

After the last round, all background processes (such as the Balancer) that had been previously disabled are now re-enabled. The reconfiguration is now complete.

Fault-Tolerance. When there is no reconfiguration, Morplus is as fault-tolerant as MongoDB. Under ongoing reconfiguration, when there is one failure, Morplus provides similar fault-tolerance as MongoDB – in a nutshell, Morplus does not lose data, but in some cases the reconfiguration may need to be restarted partially or completely.

Consider the single failure case in detail. Consider a replica set size of $rs \geq 3$. Right after isolating the first secondary in the first round, the old data configuration is still present at $(rs - 1)$ servers: current primary and identical $(rs - 2)$ idle secondaries. If the primary or an idle secondary fails, reconfiguration remains unaffected. If the currently-reconfiguring secondary fails, then the reconfiguration can be continued using one of the idle secondaries (from that replica set) instead; when the failed secondary recovers it participates in a subsequent reconfiguration round.

In a subsequent round (≥ 2), if one of the non-reconfigured replicas fails, it recovers and catches up directly with the re-configured primary. Only in the second round, if the already-reconfigured primary fails, does the entire reconfiguration need to be restarted as this server was not replicated yet. Writes between the new primary election (Round 1 Commit phase) up to its failure, before the second round completes, may be lost. This is similar to the loss of a normal MongoDB write which happens when a primary fails before replicating the data to the secondary. The vulnerability window is longer in Morplus, although this can be reduced by using a backup Morplus server.

With multiple failures, Morplus is fault-tolerant under some combinations. For instance, if all replica sets have at least one new-configuration replica left, or if all replica sets have at least one old-configuration replica left. In the former case, failed replicas can catch up. In the latter case, reconfiguration can be restarted using the surviving replicas.

3 ALGORITHMS FOR EFFICIENT SHARD KEY RE-CONFIGURATIONS

A reconfiguration operation entails the data present in shards across multiple servers to be resharded. The new shards need to be placed at the servers in such a way as to: 1) reduce the total network transfer volume during reconfiguration, and 2) achieve load balance. This section presents optimal algorithms for this planning problem.

We present two algorithms for placement of the new chunks in the cluster. Our first algorithm is greedy and is optimal in the total network transfer volume. However, it may create bottlenecks by clustering many new chunks at a few servers. Our second algorithm, based on bipartite matching, is optimal in network transfer volume among all those strategies that ensure load balance.

3.1 Greedy Assignment

The greedy approach considers each new chunk independently. For each new chunk NC_i , the approach evaluates all the N servers. For each server S_j , it calculates the number of data items W_{NC_i, S_j} of chunk NC_i that are already present in old chunks at server S_j . The approach then allocates each new chunk NC_i to that server S_j which has the maximum value of W_{NC_i, S_j} , i.e., $\text{argmax}_{S_j} (W_{NC_i, S_j})$. As chunks are considered independently, the algorithm produces the same output irrespective of the order in which chunks are considered by it.

The calculation of W_{NC_i, S_j} values can be performed in parallel at each server S_j , after servers are made aware of the new chunk ranges. A centralized server collects all the W_{NC_i, S_j} values, runs the greedy algorithm, and informs the servers of the allocation decisions.

Lemma 1. The greedy algorithm is optimal in total network transfer volume.

Proof. The proof is by contradiction. Consider an alternative optimal strategy A that assigns at least one new chunk NC_i to a server S_k different from $S' = \text{argmax}_{S_j} (W_{NC_i, S_j})$, such that $W_{NC_i, S'} > W_{NC_i, S_k}$ – if there is no such NC_i , then A produces the same total network transfer volume as the greedy approach. By instead changing A so that NC_i is re-assigned to S' , one can achieve a reconfiguration that has a lower network transfer volume than A , a contradiction. \square

For each of the m new chunks, this algorithm iterates through all the N servers. Thus its complexity is $O(m.N)$, linear in the number of new chunks and cluster size.

To illustrate the greedy scheme in action, Fig. 2 provides two examples for the shard key change operation. In each example, the database has 3 old chunks $OC_1 - OC_3$ each containing 3 data items. For each data item, we show the old shard key K_o and the new shard key K_n (both in the ranges 1-9). The new configuration splits the new key range evenly across 3 chunks shown as $NC_1 - NC_3$.

In Fig. 2a, the old chunks are spread evenly across servers $S_1 - S_3$. The edge weights in the bipartite graph show the number of data items of NC_i that are local at S_j , i.e., W_{NC_i, S_j} values. Thick lines show the greedy assignment.

However, the greedy approach may produce an unbalanced chunk assignment for skewed bipartite graphs, as in Fig. 2b. While the greedy approach minimizes network transfer volume, it assigns new chunks NC_2 and NC_3 to server S_1 , while leaving server S_3 empty.

3.2 Load Balance via Bipartite Matching

Load balancing chunks across servers is important for several reasons: i) it improves read/write latencies for clients by spreading data and queries across more servers; ii) it reduces read/write bottlenecks; iii) it reduces the tail of the reconfiguration time, by preventing allocation of too many chunks to any one server.




Fig. 2: Greedy and Hungarian strategy for shard key change using: (a) Balanced, (b) Unbalanced old chunk configuration. $S_1 - S_3$ represent servers. $OC_1 - OC_3$ and $NC_1 - NC_3$ are old and new chunks respectively. K_o and K_n are old and new shard keys respectively. Edges are annotated with W_{NC_i, S_j} weights.

Our second strategy achieves load balance by capping the number of new chunks allocated to each server. With m new chunks, this per-server cap is $\lceil m/N \rceil$ chunks. We then create a bipartite graph with two sets of vertices – top and bottom. The top set consists of $\lceil m/N \rceil$ vertices for each of the N servers in the system; denote the vertices for server S_j as $S_j^1 - S_j^{\lceil m/N \rceil}$. The bottom set of vertices consist of the new chunks. All edges between a top vertex S_j^k and a bottom vertex NC_i have an edge cost equal to $|NC_i| - W_{NC_i, S_j}$ i.e., the number of data items that will move to server S_j if new chunk NC_i were allocated to it.

Assigning new chunks to servers in order to minimize data transfer volume now becomes a bipartite matching problem. Thus, we find the minimum weight matching by using the classical Hungarian algorithm [10]. The complexity of this algorithm is $O((N.V + m).N.V.m)$ where $V = \lceil m/N \rceil$ chunks. This reduces to $O(m^3)$. The greedy strategy of Section 3.1 becomes a special case of this algorithm with $V = m$.

Lemma 2. Among all load-balanced strategies that assign at most $V = \lceil m/N \rceil$ new chunks to any server, the Hungarian algorithm is optimal in total network transfer volume.

Proof. The proof follows from the optimality of the Hungarian algorithm [10]. \square

Fig. 2b shows the outcome of the bipartite matching algorithm using dotted lines in the graph. While it incurs the same overall cost as the greedy approach, it additionally provides the benefit of a load-balanced new configuration, where each server is allocated exactly one new chunk.

While we focus on the shard key change, this technique can also be used for other reconfigurations like changing shard size, or cluster scale out and scale in. The bipartite graph would be drawn

appropriately (depending on the reconfiguration operation), and the same matching algorithm used. For purpose of concreteness, the rest of the paper focuses on shard key change.

Finally, although we have used datasize (number of key-value pairs) as the main cost metric. Instead we could use traffic to key-value pairs as the cost metric, and derive edge weights in the bipartite graph (Fig. 2) from these traffic estimates. Hungarian approach on this new graph would balance out traffic load, while trading off optimality – further exploration of this variant is beyond our scope in this paper.

4 NETWORK-AWARENESS

In this section, we describe how we augment the design of Section 2 in order to handle two important concerns: awareness to the topology of a datacenter, and geo-distributed settings.

4.1 Awareness to Datacenter Topology

Datacenters use a wide variety of topologies, the most popular being hierarchical, e.g., a typical two-level topology consists of a core switch and multiple rack switches. Others that are commonly used in practice include fat-trees [17], CLOS [36], and butterfly [35].

Our first-cut data migration strategy discussed in Section 2 was *chunk-based*: it assigned as many sockets (TCP streams) to a new chunk C at its destination server as there are source servers for C i.e., it assign one TCP stream per server pair. Using multiple TCP streams per server pair has been shown to better utilize the available network bandwidth [25]. Further, the chunk-based approach also results in *stragglers* in the execution phase as shown




Fig. 3: Execution phase CDF for chunk-based strategy on Amazon 500 MB database in tree network topology with 9 mongod servers spread evenly across 3 racks.

in Figure 3. Particularly, we observe that 60% of the chunks finish quickly, followed by a 40% cluster of chunks that finish late.

To address these two issues, we propose a weighted fair sharing (WFS) scheme that takes both data transfer size and network latency into account. Consider a pair of servers i and j , where i is sending some data to j during the reconfiguration. Let $D_{i,j}$ denote the total amount of data that i needs to transfer to j , and $L_{i,j}$ denote the latency in the shortest network path from i to j . Then, we set $X_{i,j}$, the weight for the flow from server i to j , as follows:

$$X_{i,j} \propto D_{i,j} \times L_{i,j}$$

In our implementation, the weights determine the number of sockets that we assign to each flow. We assign each destination server j a total number of sockets $X_j = K \times \frac{\sum_i D_{i,j}}{\sum_{i,j} D_{i,j}}$, where K is the total number of sockets throughout the system. Thereafter each destination server j assigns each source server i a number of sockets, $X_{i,j} = X_j \times \frac{C_{i,j}}{\sum_i C_{i,j}}$.

However, $X_{i,j}$ may be different from the number of new chunks that j needs to fetch from i . If $X_{i,j}$ is larger, we treat each new chunk as a data slice, and iteratively split the largest slice into smaller slices until $X_{i,j}$ equals the total number of slices. Similarly, if $X_{i,j}$ is smaller, we use iterative merging of the smallest slices. Finally, each slice is assigned a socket for data transfer. Splitting or merging slices is only for the purpose of socket assignment and to speed up data transfer; it does not affect the final chunk configuration which was computed in the Prepare phase.

Our approach above could have used estimates of available bandwidth instead of latency estimates. We chose the latter because: i) they can be measured with a lower overhead, ii) they are more predictable over time, and iii) they are correlated to the effective bandwidth.

4.2 Geo-Distributed Settings

So far, Morpheus has assumed that all its servers reside in one datacenter. However, typical NoSQL configurations split servers across geo-distributed datacenters for fault-tolerance. Naively using the Morpheus system would result in bulk transfers across the wide-area network and prolong reconfiguration time.

To address this, we localize each stage of the data transfer to occur within a datacenter. We leverage MongoDB’s server tags [12] to tag each replica set member with its datacenter

identifier. Morpheus then uses this information to select replicas, which are to be reconfigured together in each given round, in such a way that they reside within the same datacenter. If wide-area transfers cannot be eliminated at all, Morpheus warns the database admin.

One of MongoDB’s invariants for partition-tolerance requires each replica set to have a voting majority at some datacenter [12]. In a three-member replica set, two members (primary and secondary-1) must be at one site while the third member (secondary-2) could be at a different site. Morpheus obeys this requirement by selecting that secondary for the first round which is co-located with the current primary. In the above example, Morpheus would select the secondary-1 replicas for the first round of reconfiguration. In this way, the invariant stays true even after the leader election in the Commit phase.

5 EVALUATION

Our experiments are designed to answer the following questions:

- How much does Morpheus affect read and write operations during reconfiguration?
- For shard key change, how do the Greedy and Hungarian algorithms of Section 3 compare?
- How does Morpheus scale with data size, operation injection rate, and cluster size?
- How much benefit can we expect to get from the network-aware (datacenter topology and geo-distributed) strategies?

5.1 Setup

Data Set. We use the dataset of Amazon reviews as our default collection [38]. Each data item has 10 fields. We choose *productID* as the old shard key, *userID* as the new shard key, while update operations use these two fields and a *price* field. Our default database size is 1 GB (we later show scalability with data size).

Cluster. The default Morpheus cluster uses 10 machines. These consist of one mongos (front-end), and 3 replica sets, each containing a primary and two secondaries. There are 3 config servers, each co-located on a physical machine with a replica set primary – this is an allowed MongoDB installation. All physical machines are d710 Emulab nodes [5] with a 2.4 GHz processor, 4 cores, 12 GB RAM, 2 hard disks of 250 GB and 500 GB, 64 bit CentOS 5.5, and connected to a 100 Mbps LAN switch.

Workload Generator. We implemented a custom workload generator that injects YCSB-like workloads via MongoDB’s *pymongo* interface. Our default injection rate is 100 ops/s with 40% reads, 40% updates, and 20% inserts. To model realistic key access patterns, we select keys for each operation via one of three YCSB-like [27] distributions: 1) Uniform (default), 2) Zipf, and 3) Latest. For Zipf and Latest distributions we employ a shape parameter $\alpha = 1.5$. The workload generator runs on a dedicated pc3000 node in Emulab running a 3GHz processor, 2GB RAM, two 146 GB SCSI disks, 64 bit Ubuntu 12.04 LTS.

Morpheus default settings. Morpheus was implemented in about 4000 lines of C++ code. The code is publicly available at <http://dprg.cs.uiuc.edu/downloads>. A demo of Morpheus can be found at <https://youtu.be/0rO2oQyyg0o>. For the evaluation, each plotted datapoint is an average of at least 3 experimental trials, shown along with standard deviation bars.




Fig. 4: Read Latency for: (a) Read only operations (no writes), and three read-write workloads modeled after YCSB: (b) Uniform, (c) Zipf, and (d) Latest. Times shown are hh:mm:ss. Failed reads are shown as negative latencies. Annotated “Primary Change” point marks the start of the leader election protocol in the first round.

5.2 Read Latency

Fig. 4 shows the timelines for four different workloads during the reconfiguration, lasting between 6.5 minutes to 8 minutes. The figure depicts the read latencies for the reconfigured database table (collection), with failed reads shown as negative latencies. We found that read latencies for collections not being reconfigured were not affected and we do not plot these.

Fig. 4a shows the read latencies when there are no writes (Uniform read workload). We observe unavailability for a few seconds (from time $t=18:28:21$ to $t=18:28:29$) during the Commit phase when the primaries are being changed. This unavailability lasts only about 2% of the total reconfiguration time. After the change, read latencies spike slightly for a few reads but then settle down. Figs. 4b to 4d plot the YCSB-like read-write workloads. We observe similar behavior as Fig. 4a.

	Read	Write
Read Only	99.9	-
Uniform	99.9	98.5
Latest	97.2	96.8
Zipf	99.9	98.3

TABLE 1: Percentage of Reads and Writes that Succeed under Reconfiguration. (Figs. 4 and 6)

Fig. 4d indicates that the Latest workload incurs a lot more failed reads. This is because the keys that are being inserted and updated are the ones more likely to be read. However, since some of the insertions fail due to the write throttles at various points during the reconfiguration process ($t=14:11:20$ to $t=14:11:30$,

$t=14:13:15$ to $t=14:13:20$, $t=14:15:00$ to $t=14:15:05$ from Fig. 6c), this causes subsequent reads to also fail. However, these lost reads only account for 2.8% of the total reads during the reconfiguration time – in particular, Table 1 (middle column) shows that 97.2% of the reads succeed under the Latest workload. The availability numbers are higher at three-9’s for Uniform and Zipf workload, and these are comparable to the case when there are no insertions. We conclude that unless there is temporal and spatial (key-wise) correlation between writes and reads (i.e., Latest workloads), the read latency is not affected much by concurrent writes. When there is correlation, Morpheus mildly reduces the offered availability.

To flesh this out further, we plot in Fig. 5 the CDF of read latencies for these four settings, and when there is no reconfiguration (Uniform workload). Notice that the horizontal axis is logarithmic scale. We only consider latencies for successful reads. We observe that the 96th percentile latencies for all workloads are within a range of 2 ms. The median (50th percentile) latency for No Reconfiguration is 1.4 ms, and this median holds for both the Read only (No Write) and Uniform workloads. The medians for Zipf and Latest workloads are lower at 0.95 ms. This lowered latency is due to two reasons: caching at the mongod servers for the frequently-accessed keys, and in the case of Latest the lower percentage of successful reads.

We conclude that under reconfiguration, the read availability provided by Morpheus is high (two to three 9’s of availability), while the latencies of successful read operations do not degrade compared to the scenario when there is no reconfiguration in




Fig. 6: Write Latency for three read-write workloads modeled after YCSB: (a) Uniform, (b) Zipf, and (c) Latest. Times shown are hh:mm:ss. Failed writes are shown as negative latencies. Annotations marked “WT” indicate the start of each write throttle phase. (d) CDF of Write Latency Distribution for no reconfiguration (No Reconf) and three under-reconfiguration workloads.




Fig. 5: CDF of Read Latency. Read latencies under no reconfiguration (No Reconf), and four under-reconfiguration workloads.

progress.

5.3 Write Latency

We next plot the data for write operations, i.e., inserts, updates and deletes.

Figs. 6a to 6c show writes in the same timelines as Figs. 4b to 4d. We observe that many of the failed writes occur during one of the write throttling periods (annotated as “WT”). Recall from Section 2.2 that there are as many write throttling periods

as the replica set size, with one throttle period at the end of each reconfiguration round.

Yet, the number of writes that fail is low: Table 1 (last column) shows that for the Uniform and Zipf workloads, fewer than 2% writes fail. The Latest workload again has a slightly higher failure rate since a key that was attempted to be written (unsuccessfully) is more likely to be attempted to be written again in the near future. Yet, the write failure rate of 3.2% is reasonably low.

To flesh this out further, the CDF of the write latencies (ignoring failed writes) is shown in Fig. 6d. The median for writes when there is no reconfiguration (Uniform workload) in progress is 1.45 ms. The Zipf and Latest workloads have a similar median latency. Uniform has a slightly higher median latency at 1.6 ms – this is because 18% of the updates experience high latencies. This is due to Greedy’s skewed chunk assignment plan (discussed and improved in Section 5.4). Greedy assigns a large percentage of new chunks to a single replica set, which thus receives most of the new write traffic. This causes MongoDB’s periodic write journal flushes⁴ to take longer. This in turn delays the new writes arriving around the journal flush timepoints. Many of the latency spikes observed in Figs. 6a to 6c arise from this journaling behavior.

We conclude that under reconfiguration, the write availability provided by Morpheus is high (close to two 9’s), while latencies of successful writes degrade only mildly compared to when there is no reconfiguration in progress.

⁴ MongoDB maintains a write-ahead (journal) log for durability which is periodically flushed to disk.

5.4 Hungarian vs. Greedy Reconfiguration




Fig. 7: Greedy (Morphus-G) vs. Hungarian (Morphus-H) Strategies for shard key change. Uncorrelated: random old and new shard key. Correlated: new shard key is reverse of old shard key.

Section 3 outlined two algorithms for the shard key change reconfiguration – Hungarian and Greedy. We implemented both these techniques into Morphus – we call these variants as Morphus-H and Morphus-G respectively. For comparison, we also implemented a random chunk assignment scheme called Morphus-R.

Fig. 7 compares these three variants under two scenarios. The uncorrelated scenario (left pair of bars) uses a synthetic 1 GB dataset where for each data item, the value for its new shard key is selected at random and uncorrelated to its old shard key’s value. The plot shows the total reconfiguration time including all the Morphus phases. Morphus-G is 15% worse than Morphus-H and Morphus-R. This is because Morphus-G ends up assigning 90% new chunks to a single replica set which results in stragglers during the Execution and Recovery phases. The underlying reason for the skewed assignment can be attributed to MongoDB’s split algorithm which we use modularly. The algorithm partitions the total data size instead of total record count. When partitioning the data using the old shard key, this results in some replica sets getting a larger number of records than others. Morphus-R performs as well as Morphus-H because by randomly assigning chunks to servers, it also achieves load balance

The correlated scenario in Fig. 7 (right pair of bars) shows the case where new shard keys have a reversed order compared to old shard keys. That is, with M data items, old shard keys are integers in the range $[1, M]$, and the new shard key for each data item is set as $= M - \text{old shard key}$. This results in data items that appeared together in chunks continuing to do so (because chunks are sorted by key). Morphus-R is 5x slower than both Morphus-G and Morphus-H. Randomly assigning chunks can lead to unnecessary data movement. In the correlated case, this effect is accentuated. Morphus-G is 31% faster than Morphus-H. This is because the total transfer volume is low anyway in Morphus-G due to the correlation, while Morphus-H additionally attempts to load-balance.

We conclude that i) the algorithms of Section 3 give an advantage over random assignment, especially when old and new keys are correlated, and ii) Morphus-H performs reasonably well in both the correlated and uncorrelated scenario and should be preferred over Morphus-G and Morphus-R.

5.5 Scalability

We explore scalability of Morphus along three axes – database size, operation injection rate, and size of cluster. These experiments use the Amazon dataset.

Database Size. Fig. 8a shows the reconfiguration time at various data sizes from 1 GB to 10 GB. There were no reads or writes injected. For clarity, the plotted data points are perturbed slightly horizontally.

Firstly, Fig. 8a shows that Morphus-H performs slightly better than Morphus-G for the real-life Amazon dataset. This is consistent with our observations in Section 5.4 since the Amazon workload is closer to the uncorrelated end of the spectrum.

Secondly, the total reconfiguration time appears to increase superlinearly beyond 5 GB. This can be attributed to two factors. First, reconfiguration time grows with the *number of chunks* – this number is also plotted, and we observe that it grows superlinearly with dataseize. This is again caused by MongoDB’s splitting code⁵. Second, we have reused MongoDB’s data transfer code, which relies on cursors (i.e., iterators), which are not the best approach for bulk transfers. We believe this can be optimized further by writing a module for bulk data transfer – yet, we reiterate that this is orthogonal to our contributions: even during the (long) data transfer time, reads and writes are still supported with several 9s of availability (Table 1). Today’s existing approach of exporting/reimporting data with the database shut down, leads to long unavailability periods – at least 30 minutes for 10 GB of data (assuming 100% bandwidth utilization). In comparison, Morphus is unavailable in the worst-case (from Table 1) for $3.2\% \times 2 \text{ hours} = 3.84 \text{ minutes}$, which is an improvement of about 10x.

Fig. 8a also illustrates that a significant fraction of the reconfiguration time is spent migrating data, and this fraction grows with increasing data size – at 10 GB, the data transfer occupies 90% of the total reconfiguration time. This indicates that Morphus’ overheads fall and will become relatively small at large data sizes.

We conclude that the reconfiguration time incurred by Morphus scales linearly with the number of chunks in the system and that the overhead of Morphus falls with increasing data size.

Operation Injection Rate. An important concern with Morphus is how fast it plays “catch up” when there are concurrent writes during the reconfiguration. Fig. 8b plots the reconfiguration time against the write rate on 1 GB of Amazon data. In both Morphus-G and Morphus-H, we observe a linear increase. More concurrent reads and writes slow down the overall reconfiguration process because of two reasons: limited bandwidth available for the reconfiguration data transfers, and a longer oplog that needs to be replayed during the Recovery Phase. However, this increase is slow and small. A 20-fold increase in operation rate from 50 ops/s to 1000 ops/s results in only a 35% increase in reconfiguration time for Morphus-G and a 16% increase for Morphus-H.

To illustrate this further, the plot shows the phase breakdown for Morphus-H. The Recovery phase grows as more operations need to be replayed. Morphus-H has only a sublinear growth in reconfiguration time. This is because of two factors. First Morphus-H balances the chunks out more than Morphus-G, and as a result the oplog replay has a shorter tail. Second, there is an overhead in Morphus-H associated with fetching the oplog via the MongoDB cursors (iterators) – at small write rates, this overhead dominates but as the write rate increases, the contribution of this

5. Our results indicate that MongoDB’s splitting algorithm may be worth revisiting.




Fig. 8: **Morphus Scalability with: (a) Data Size, also showing Morphus-H phases, (b) Operation injection rate, also showing Morphus-H phases, (c) Number of replica sets, and (d) Replica set size.**

overhead drops off. This second factor is present in Morphus-G as well, however it is offset by the unbalanced distribution of new chunks.

We conclude that Morphus catches up quickly when there are concurrent writes, and that its reconfiguration time scales linearly with write operation injection rate.

Cluster Size. We investigate cluster size scalability along two angles: number of replica sets, and replica set size. Fig. 8c shows that as the number of replica sets increases from 3 to 9 (10 to 28 servers), both Morphus-G and Morphus-H eventually become faster with scale. This is primarily because of increasing parallelism in the data transfer, while the amount of data migrating over the network grows much slower – with N replica sets, this latter quantity is approximately a fraction $\frac{N-1}{N}$ of data. While Morphus-G’s completion time is high at a medium cluster size (16 servers) due to its unbalanced assignment, Morphus-H shows a steady improvement with scale and eventually starts to plateau as expected.

Next, Fig. 8d shows the effect of increasing replica set size. We observe a linear increase for both Morphus-G and Morphus-H. This is primarily because there are as many rounds inside a reconfiguration run as there are machines in a replica set.

We conclude that Morphus scales reasonably with cluster size – in particular, an increase in number of replica sets improves its performance.

5.6 Effect of Network Awareness

Datacenter Topology-Awareness. First, Fig. 9a shows the length of the Execution phase (using a 500 MB Amazon collection)

for two hierarchical topologies, and five migration strategies. The topologies are: i) homogeneous: 9 servers distributed evenly across 3 racks, and ii) heterogeneous: 3 racks contain 6, 2, and 1 servers respectively. The switches are Emulab pc3000 nodes and all links are 100 Mbps. The inter-rack and intra-rack latencies are 2 ms and 1 ms respectively. The five strategies are: a) Fixed sharing, with one socket assigned to each destination node, b) chunk-based approach (Section 4.1), c) Orchestra [25] with $K = 21$, d) WFS with $K = 21$ (Section 4.1), and e) WFS with $K = 28$.

We observe that in the homogeneous clusters, WFS strategy with $K = 28$ is 30% faster than fixed sharing, and 20% faster than the chunk-based strategy. Compared to Orchestra which only weights flows by their data size, taking the network into account results in a 9% improvement in WFS with $K = 21$. Increasing K from 21 to 28 improves completion time in the homogeneous cluster, but causes degradation in the heterogeneous cluster. This is because a higher K results in more TCP connections, and at $K = 28$ this begins to cause congestion at the rack switch of 6 servers. Second, Fig. 9b shows that compared to Fig. 3 (from Section 4), Morphus’ network-aware WFS strategy has a shorter tail and finishes earlier. Network-awareness lowers the median chunk finish time by around 20% in both the homogeneous and heterogeneous networks.

We conclude that WFS strategy improves performance compared to existing approaches, and K should be chosen high enough but without causing congestion.

Geo-Distributed Setting. Table 2 shows the benefit of the tag-aware approach of Morphus (Section 4.2). The setup has two datacenters with 6 and 3 servers, with intra- and inter-datacenter





Fig. 9: (a) Execution Phase Migration time for five strategies: (i) Fixed Sharing (FS), (ii) Chunk-based strategy, (iii) Orchestra with $K = 21$, (iv) WFS with $K = 21$, and (v) WFS with $K = 28$. (b) CDF of total reconfiguration time in chunk-based strategy vs. WFS with $K = 28$.

latencies of 0.07 ms and 2.5 ms respectively (based on [40]) and links with 100 Mbps bandwidth. For 100 ops/s workload on 100 MB of reconfigured data, tag-aware Morpheus improves performance by over 2x when there are no operations and almost 3x when there are reads and writes concurrent with the reconfiguration.

	Without Read/Write	With Read/Write
Tag-Unaware	49.074s	64.789s
Tag-Aware	21.772s	23.923s

TABLE 2: Reconfiguration Time in the Geo-distributed setting.

5.7 Large Scale Experiment

In this experiment, we increase data and cluster size simultaneously such that the amount of data per replica set is constant. We ran this experiment on Google Cloud [6]. We used n1-standard-4 VMs each with 4 virtual CPUs and 15 GB of memory. The disk capacity was 1 GB and the VMs were running Debian 7. We generated a synthetic dataset similar to the one used in Section 5.4. Morpheus-H was used for reconfiguration with WFS migration scheme and $K =$ number of old chunks.




Fig. 10: Running Morpheus-H with WFS ($K =$ number of old chunks) for reconfiguring database of size (25GB, 50GB, 100GB) running on a cluster size (25 machines (8 replica sets * 3 + 1 mongos), 49 machines (16 replica sets) and 100 machines (33 replica sets)).

Fig. 10 shows a sublinear increase in reconfiguration time as data and cluster size increases. Note that x-axis uses log scale. In the Execution phase, all replica sets communicate among each other for migrating data. As the number of replica sets increases with cluster size, the total number of connections increases leading to network congestion. Thus, the Execution phase takes longer.

The amount of data per replica set affects reconfiguration time super-linearly. On the contrary, cluster size has a sublinear impact. In this experiment, the latter dominates as the amount of data per replica set is constant.

6 RELATED WORK

Research in distributed databases has focused on query optimization and load-balancing [23], and orthogonally on using group communication for online reconfiguration [34], however, they do not solve the core algorithmic problems for efficient reconfiguration. Online schema change was targeted in [41], but the resultant availabilities were lower than those provided by Morpheus. In a parallel work, Elmore et. al. [32] have looked into the reconfiguration problem for a partitioned main memory database like H-Store. Data placement in parallel databases have used hash-based and range-based partitioning [28], [39], but they do not target optimality for reconfiguration.

The problem of live migration has been looked into in the context of databases. Albatross [31], Zephyr [33] and ShuttleDB [21] addresses live migration in multi-tenant transactional databases. Albatross and ShuttleDB uses iterative operation replay like Morpheus, while Zephyr routes updates based on current data locations. Data migration in these systems happen between two different sets of servers while Morpheus achieves this inside the same replica sets. Also, they do not propose optimal solutions for any reconfiguration operation. Opportunistic lazy migration explored in Relational Cloud [30] entails longer completion times. Tuba [19] looked into the problem of migration in a geo-replicated setting. They avoided write thottle by having multiple masters at the same time. MongoDB does not support multiple masters in a single replica set, which dictated Morpheus's current design.

Morpheus' techniques naturally bear some similarities with live VM migration. Pre-copy techniques migrate a VM without stopping the OS, and if this fails then the OS is stopped [26]. Like pre-copy, Morpheus also replays operations that occurred during

the migration. Pre-copy systems also use write throttling [22]. Pre-copy has been used in database migration [20].

For network flow scheduling, Chowdhury et al [25] proposed a weighted flow scheduling which allocates multiple TCP connections to each flow to minimize migration time. Our WFS approach improves their approach by additionally considering network latencies. Morphus' performance is likely to improve further if we also consider bandwidth. Hedera [18] also provides a dynamic flow scheduling algorithm for multi-rooted network topology. Even though these techniques may improve reconfiguration time, Morphus' approach is end-to-end and is less likely to disrupt normal reads and writes which use the same network links.

7 SUMMARY

This paper described optimal and load-balanced algorithms for online reconfiguration operation, and the Morphus system integrated into MongoDB. Our experiments showed that Morphus supports fast reconfigurations such as shard key change, while only mildly affecting the availability and latencies for read and write operations. Morphus scales well with data size, operation injection rate, and cluster size.

REFERENCES

- [1] Altering Cassandra column family primary key. <http://stackoverflow.com/questions/18421668/alter-cassandra-column-family-primary-key-using-cassandra-cli-or-cql>. visited on 2015-1-5.
- [2] Change shard key mongodb faq. <http://docs.mongodb.org/manual/faq/sharding/#can-i-change-the-shard-key-after-sharding-a-collection>. visited on 2015-1-5.
- [3] Command to change shard key of a collection. <https://jira.mongodb.org/browse/SERVER-4000>. visited on 2015-1-5.
- [4] CouchDB. <http://couchdb.apache.org>. visited on 2015-1-5.
- [5] Emulab. <https://wiki.emulab.net/wiki/d710>. visited on 2015-1-5.
- [6] Google Cloud. <https://cloud.google.com/>. visited on 2015-1-5.
- [7] The great primary key debate. <http://www.techrepublic.com/article/the-great-primary-key-debate/>. visited on 2015-1-5.
- [8] HBase. <https://hbase.apache.org>. visited on 2015-1-5.
- [9] How to change the Shard Key. <http://stackoverflow.com/questions/6622635/how-to-change-the-shard-key>. visited on 2015-1-5.
- [10] Hungarian algorithm. http://en.wikipedia.org/wiki/Hungarian_algorithm. visited on 2015-1-5.
- [11] MongoDB. <http://www.mongodb.org>. visited on 2015-1-5.
- [12] Mongodb manual for geo-distributed deployment (2.4.9). <http://docs.mongodb.org/manual/tutorial/deploy-geographically-distributed-replica-set/>. visited on 2015-1-5.
- [13] RethinkDB. <http://rethinkdb.com/>. visited on 2015-1-5.
- [14] SMG Research Reveals DBA Tools Not Effective for Managing Database Change. <http://www.datical.com/news/research-reveals-dba-tools-not-effective-for-database-change/>. visited on 2015-04-11.
- [15] Troubles With Sharding - What Can We Learn From The Foursquare Incident? <http://highscalability.com/blog/2010/10/15/troubles-with-sharding-what-can-we-learn-from-the-foursquare.html>. visited on 2015-04-11.
- [16] NoSQL market forecast 2015-2020, Market Research Media. <http://www.marketresearchmedia.com/?p=568>, 2012. visited on 2015-1-5.
- [17] M. Al-Fares, S. Loukissas, and A. Vahdat. A scalable, commodity data center network architecture. In *Proceedings of 2008 the ACM Conference on Special Interest Group on Data Communication*, SIGCOMM '08, pages 63–74, New York, NY, USA, 2008. ACM.
- [18] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. Hedera: Dynamic flow scheduling for data center networks. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI '10, pages 19–19, Berkeley, CA, USA, 2010. USENIX Association.
- [19] M. S. Ardekani and D. B. Terry. A self-configurable geo-replicated cloud storage system. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*, OSDI '14, pages 367–381, Berkeley, CA, USA, 2014. USENIX Association.
- [20] S. Barker, Y. Chi, H. J. Moon, H. Hacigümüş, and P. Shenoy. “Cut me some slack”: Latency-aware live migration for databases. In *Proceedings of the 15th International Conference on Extending Database Technology*, EDBT '12, pages 432–443, New York, NY, USA, 2012. ACM.
- [21] S. K. Barker, Y. Chi, H. Hacigümüş, P. J. Shenoy, and E. Cecchet. Shuttledb: Database-aware elasticity in the cloud. In *11th International Conference on Autonomic Computing*, ICAC '14, Philadelphia, PA, USA, June 18–20, 2014., pages 33–43, 2014.
- [22] R. Bradford, E. Kotsovinos, A. Feldmann, and H. Schiöberg. Live wide-area migration of virtual machines including local persistent state. In *Proceedings of the 3rd International Conference on Virtual Execution Environments*, VEE '07, pages 169–179, New York, NY, USA, 2007. ACM.
- [23] M. J. Carey and H. Lu. Load balancing in a locally distributed DB system. In *Proceedings of the 1986 ACM International Conference on Special Interest Group on Management of Data*, SIGMOD '86, pages 108–119, New York, NY, USA, 1986. ACM.
- [24] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. In *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation - Volume 7*, OSDI '06, pages 15–15, Berkeley, CA, USA, 2006. USENIX Association.
- [25] M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica. Managing data transfers in computer clusters with Orchestra. In *Proceedings of the ACM Special Interest Group on Data Communication 2011 Conference*, SIGCOMM '11, pages 98–109, New York, NY, USA, 2011. ACM.
- [26] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live migration of virtual machines. In *Proceedings of the 2nd Conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 273–286, Berkeley, CA, USA, 2005. USENIX Association.
- [27] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with YCSB. In *Proceedings of the 1st ACM Symposium on Cloud Computing*, SoCC '10, pages 143–154, New York, NY, USA, 2010. ACM.
- [28] G. Copeland, W. Alexander, E. Boughter, and T. Keller. Data placement in Bubba. In *Proceedings of the 1988 ACM International Conference on Special Interest Group on Management*, SIGMOD '88, pages 99–108, New York, NY, USA, 1988. ACM.
- [29] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. 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. *ACM Trans. Comput. Syst.*, 31(3):8:1–8:22, Aug. 2013.
- [30] C. Curino, E. P. C. Jones, R. A. Popa, N. Malviya, E. Wu, S. Madden, H. Balakrishnan, and N. Zeldovich. Relational cloud: a database service for the cloud. In *Proceedings of 5th Biennial Conference on Innovative Data Systems Research*, pages 235–240, 2011.
- [31] S. Das, S. Nishimura, D. Agrawal, and A. El Abbadi. Albatross: lightweight elasticity in shared storage databases for the cloud using live data migration. In *Proceedings of the Very Large Database Endowment*, volume 4, pages 494–505. VLDB Endowment, May 2011.
- [32] A. Elmore, V. Arora, R. Taft, A. Pavlo, D. Agrawal, and A. E. Abbadi. Squall: Fine-grained live reconfiguration for partitioned main memory databases. In *Proceedings of the 2015 ACM International Conference on Special Interest Group on Management of Data*, SIGMOD '15, 2015.
- [33] A. J. Elmore, S. Das, D. Agrawal, and A. El Abbadi. Zephyr: Live migration in shared nothing databases for elastic cloud platforms. In *Proceedings of the 2011 ACM International Conference on Special Interest Group on Management of Data*, SIGMOD '11, pages 301–312, New York, NY, USA, 2011. ACM.
- [34] B. Kemme, A. Bartoli, and O. Babaoglu. Online reconfiguration in replicated databases based on group communication. In *Proceedings of the 2001 International Conference on Dependable Systems and Networks (Formerly: FTCS)*, DSN '01, pages 117–130, Washington, DC, USA, 2001. IEEE Computer Society.
- [35] J. Kim, W. J. Dally, and D. Abts. Flattened butterfly: A cost-efficient topology for high-radix networks. In *Proceedings of the 34th Annual International Symposium on Computer Architecture*, ISCA '07, pages 126–137, New York, NY, USA, 2007. ACM.
- [36] J. Kim, W. J. Dally, and D. Abts. Efficient topologies for large-scale cluster networks. In *Proceedings of the 2010 Conference on Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference (OFC/NFOEC)*, pages 1–3. IEEE, 2010.

- [37] A. Lakshman and P. Malik. Cassandra: A decentralized structured storage system. *ACM SIGOPS Operating Systems Review*, 44(2):35–40, Apr. 2010.
- [38] J. McAuley and J. Leskovec. Hidden factors and hidden topics: Understanding rating dimensions with review text. In *Proceedings of the 7th ACM Conference on Recommender Systems, RecSys '13*, pages 165–172, New York, NY, USA, 2013. ACM.
- [39] M. Mehta and D. J. DeWitt. Data placement in shared-nothing parallel database systems. *The Very Large Database Journal*, 6(1):53–72, Feb. 1997.
- [40] G. Pang. Latencies gone wild!, AMPLab - UC Berkeley. <https://amplab.cs.berkeley.edu/2011/10/20/latencies-gone-wild/>. visited on 2015-1-5.
- [41] I. Rae, E. Rollins, J. Shute, S. Sodhi, and R. Vingralek. Online, asynchronous schema change in F1. In *Proceedings of the Very Large Database Endowment*, volume 6, pages 1045–1056. VLDB Endowment, Aug. 2013.
- [42] D. Sjøberg. Quantifying schema evolution. *Information and Software Technology*, 35(1):35–44, 1993.



Mainak Ghosh is Ph.D. candidate in Computer Science of University of Illinois, Urbana Champaign. He received his masters from IIT Kharapur in 2009 and bachelors from Jadavpur University in 2007. His research interests span NoSQL databases, and real-time stream processing systems. He has previously worked at Adobe (2009-2012) and interned at IBM Research and Microsoft Research.



Wenting Wang is a Ph.D. candidate in Computer Science Department of University of Illinois, Urbana Champaign. She received her masters degree in Shanghai Jiao Tong University and her bachelors degree in Nanjing University, China. Her research interests include NoSQL database and real-time stream processing systems. She is also a software engineer in Yelp.



Gopalakrishna Holla is a Software Engineer in the Distributed Data Systems group at LinkedIn. He works on Ambry, LinkedIn's geo-distributed blob storage system. In addition to dealing with the classical distributed systems problems of consistency and replication, he is interested in building systems that can efficiently store and serve large media. For leisure, he likes to play soccer, swim or go rock climbing.



Indranil Gupta (Indy) is an Associate Professor of Computer Science at the University of Illinois at Urbana-Champaign. His Distributed Protocols Research Group (<http://dprg.cs.uiuc.edu>) works on large-scale distributed systems, with a recent focus on cloud and big data systems. Indranil is recipient of the the NSF CAREER award in 2005, and Best Paper Awards (IEEE ICAC 2015, BigMine 2012). He has (co-)chaired leading distributed computing conferences (ICDCS 2016 Cloud Computing and Datacenters Track, IEEE P2P 2014, ACM/IFIP/Usenix Middleware 2010, IEEE SASO 2010, StoDiS 2005, ACM PODC 2007 General Chair). Previously, Indranil received his PhD from Cornell University in 2004 and Bachelors degree from IIT Madras in 1998. He worked at Google (2011-12) as a full-time employee/Visiting Scientist in Mountain View, and previously at IBM Research and Microsoft Research.