

# Transactional Partitioning: A New Abstraction for Main-Memory Databases

Vivek Shah  
(supervised by Marcos Vaz Salles)  
University of Copenhagen  
Copenhagen, Denmark  
bonii@di.ku.dk

## ABSTRACT

The growth in variety and volume of OLTP (Online Transaction Processing) applications poses a challenge to OLTP systems to meet performance and cost demands in the existing hardware landscape. These applications are highly interactive (latency sensitive) and require update consistency. They target commodity hardware for deployment and demand scalability in throughput with increasing clients and data. Currently, OLTP systems used by these applications provide trade-offs in performance and ease of development over a variety of applications. In order to bridge the gap between performance and ease of development, we propose an intuitive, high-level programming model which allows OLTP applications to be modeled as a cluster of application logic units. By extending transactions guaranteeing full ACID semantics to provide the proposed model, we maintain ease of application development. The model allows the application developer to reason about program performance, and to influence it without the involvement of OLTP system designers (database designers) and/or DBAs. As a result, the database designer is free to focus on efficient running of programs to ensure optimal cluster resource utilization.

## 1. INTRODUCTION

OLTP applications have grown in volume, variety and performance demands [20]. In addition to the classic banking, reservation, and order entry systems, novel applications include massively multiplayer online worlds (MMOW) or virtual worlds, massively multiplayer online role playing games (MMORPG), financial applications (e.g., online trading), telecommunication applications and information visualizations. These applications demand scalability in throughput with growing data and clients. They exhibit high interactivity (latency sensitive) while maintaining consistency on updates. Furthermore, they typically target commodity hardware for deployment. We will refer to them as HICCUP (Highly-Interactive Commodity-Hardware Consistent-on-Update) applications.

Stonebraker et al. have shown that OLTP systems can be deployed in main-memory using shared-nothing cluster of machines at modest cost [21]. This leads to order-of-magnitude gains in performance over classical relational database management systems (RDBMS). With the reduction in prices of commodity hardware and growth of cloud computing [1], a cluster-based infrastructure consisting of shared-memory commodity hardware nodes increasingly provides flexible main-memory deployment options to meet HICCUP application requirements.

In order to build HICCUP applications for cluster-based architectures, there are currently two dominant programming models exposed to the application developer. The first approach abstracts the distributed architecture into a uniform shared-memory space, which is then presented using a unified data model, e.g., the relational model under strong consistency semantics (ACID transactions). The second approach abstracts the distributed architecture into a low-level distributed storage-oriented model, e.g., key value stores with no/loose consistency semantics.

Although the first approach provides ease of development, performance varies with the variety of applications. This is because of the sensitivity of the approach to data and code partitioning to reduce the impact of distributed transactions [11]. In the second approach, performance is controlled by the application developer, but development is extremely hard due to no/loose consistency semantics and the low-level, storage-oriented programming model. There are in-between approaches that partition the unified data model and provide consistency semantics within partitions [2, 19]. However, in these systems, the clients must specify these partitions and handle consistency across partitions. This increases application complexity, especially for HICCUP applications which cannot always be architected to eliminate inter-partition accesses.

We want to address this gap by exploring a middle ground between the two extreme approaches. Strong consistency semantics (ACID transactions) provide ease of application development. Sacrificing consistency semantics is not enough to guarantee OLTP application performance [10]. We want to build a transactional system maintaining global consistency semantics in the cluster to enable ease of development while guaranteeing optimal cluster performance. In order to do so, we propose to extend the transactional abstraction with a logical distribution abstraction. This enables the application developers to reason about the performance of their application in terms of distributed transactional logic units and thus write efficient programs.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing [info@vldb.org](mailto:info@vldb.org).  
*Proceedings of the VLDB 2014 PhD Workshop.*

While the logical distribution abstraction hides the low-level distributed architecture, it allows application developers to understand program performance behavior in a cluster setup and to influence it. Database designers can then focus on running the programs efficiently for optimal cluster resource utilization. This approach ensures that the application developers can participate in improving the performance of the application. The approach also reduces the criticality of data and code partitioning. As a result, it removes database designers from the critical path of performance improvement over a variety of applications by separating concerns of writing efficient programs (application developers) and running them efficiently (database designers). Global consistency guarantees in the cluster free the developer from worrying about consistency issues, thus reducing program complexity and development costs.

The remainder of the paper is organized as follows. In Section 2, we provide a background of the project by laying down the design requirements for OLTP systems and reviewing the state of the art with respect to these requirements. In Section 3, we describe our approach in greater detail with an example OLTP application (the TPC-C <sup>1</sup> benchmark). Finally, in Section 4, we outline the challenges and roadmap over the course of the project with some initial thoughts on how to tackle them.

## 2. BACKGROUND

In this section, we provide a background of the Ph.D. project by first isolating the OLTP design requirements in the current hardware and software landscape and then reviewing the state of the art with respect to these design requirements.

### 2.1 OLTP Design Requirements

1. In order to meet the requirements of a growing variety of HICCUP applications, OLTP systems must expose a programming model which allows application developers to reason about program performance in a cluster-based setup and to improve the performance of their architected applications (programs).
2. Lack of global consistency semantics and a low-level distributed storage-oriented programming model make application development extremely expensive, error-prone and limit the adoptability of such a system [22]. OLTP systems must provide a high-level, intuitive programming model with strong global consistency semantics over the distributed cluster setup to ease application development.
3. To take advantage of advances in commodity hardware and the growth of cloud computing [1], OLTP systems must target a hardware ecosystem of shared-memory nodes in a shared-nothing setup, allowing flexible deployments varying in performance and cost. OLTP systems must also be designed to utilize the multi-processor, multi-core processor and multi-tiered cache architecture in individual cluster nodes to ensure optimal hardware performance [13].

---

<sup>1</sup>[http://www.tpc.org/tpcc/spec/tpcc\\_current.pdf](http://www.tpc.org/tpcc/spec/tpcc_current.pdf)

## 2.2 State of the Art

In recent times there have been numerous implementations of systems that can be used by HICCUP applications. There has been an explosion in the implementation of key-value stores, which provide varying performance, scalability and availability characteristics. Amazon Dynamo [8], Google BigTable [3], Yahoo! PNUTS [5], Amazon HBase and Cassandra [15], Amazon S3, Google Cloud Storage, Windows Azure Storage are some examples of distributed key-value stores. These systems target commodity hardware in a cluster-based architecture and provide high availability and scalability at the cost of lower consistency semantics and a low-level storage-oriented programming model. In order to build HICCUP applications using these systems, applications must implement multi-key transactions and manage the low-level key-value based infrastructure, which is extremely complex and error prone, thus limiting ease of development [22].

In order to fix these problems, transactional capabilities were introduced in distributed data stores, such as Google Spanner [6] and Megastore [2], or exposed using client libraries [9, 18]. Megastore [2] provides full ACID semantics but only within partitions of the data, which limits the ease of application development and flexible deployments. Warp [9] and Percolator [18] provide multi-key transactions using client libraries but their storage-oriented programming model and client-based execution make it difficult to reason about application performance for varying cluster-based deployments.

On the other end, classical RDBMS provide a high-level uniform shared-memory abstraction, which makes application development extremely easy. However, these systems face performance scalability issues. In order to meet these challenges, Shore-MT [13] and DORA [16] target shared-memory architectures for high performance. Shore-MT optimizes a classical RDBMS engine by reducing or removing contention points for multi-core and multi-processor hardware. DORA proposes a novel methodology to partition data amongst the processing threads and migrates transactions between threads based on the partitioning. Although both these systems ease application development, reasoning about program performance remains extremely hard. These systems are also not suited for cluster-based deployment unless a shared-memory layer is built underneath them.

H-Store was one of the first systems to target cluster-based deployments for OLTP applications [14]. In order to provide a high-level relational abstraction, H-Store is reliant on inferring optimal data partitioning from application logic to minimize the impact of distributed transactions [7]. This makes it extremely difficult for application developers to reason about program performance over a variety of applications and increases the performance reliance of the system on automatic data [17] and code partitioning [4], or database administrators (DBAs). This brings database designers and/or DBAs in the critical path of performance improvements of varying OLTP applications and ignores application developers as a valuable resource.

Our approach is to provide a high-level, intuitive programming model over a cluster-based architecture where the uniform shared-memory space is partitioned by application logic. This partitioning is done by the application developers, which allows them to reason about the performance of their programs and to improve them. By extending

transactions to provide this abstraction with global ACID semantics, our approach maintains ease of application development.

### 3. APPROACH

In this section, we explain our approach in detail with an example HICCUP application (the TPC-C benchmark) and demonstrate how it meets the design requirements outlined in Section 2.

#### 3.1 Logical Partitions

In a cluster-based architecture, the programming model must provide an abstraction to balance ease of development and utilization of cluster resources. We want to explore the middle ground between the two current dominant approaches by allowing programmers to reason about the performance of their programs in the distributed setup while maintaining ease of use. We do so by introducing the notion of logical partitions. A logical partition forms the unit of logical code isolation. A programmer architects his application to consist of multiple application logic units, which communicate with each other, and establishes control flow dependencies between them. An application logic unit must be associated with a logical partition by the programmer. Communication between the application logic units is expensive, which encourages the developers to think about the performance of the constructed program and to make it more efficient. In order to provide ease of development, ACID semantics are maintained globally across logical partitions via transactions, which form the unit of code construction in this model.

In our model, a program (transaction) begins execution on a logical partition. During its execution the program can only access data on that logical partition. If it needs to access data on another logical partition, it must explicitly do so by invoking a program on another logical partition which is executed as a subtransaction [25, p. 253]. When a transaction or subtransaction is invoked, the logical partition where it must be executed must be specified as well. Program code is pre-compiled on each logical partition and executed within the database process for performance.

In contrast to partitioning data, our model forces the application developer to partition his application logic. The programmer must construct his application in terms of distributed logical partitions by associating application logic to them. Data partitioning is a consequence of the code partitioning formulated by the application developer. Existing OLTP systems either provide the view of a single data partition or fragmented data partitions, but not partitioning in terms of application logic. This increases the reliance of these systems on optimal code and data partition layout, thus making it hard for application developers to reason about the performance of a variety of applications using these systems.

In summary, our model allows the application developer to reason about program performance in terms of distributed application logic units by exposing coordination costs between logical partitions. The database designer can focus on providing efficient mapping of logical partitions to physical partitions (machines) for optimal cluster resource utilization. Our model allows ease of construction and deployment of different HICCUP applications across various cluster configurations.

```

txn new_order(w_id, d_id, c_id, order) {
  //Get the order id from the district and insert the
  //next order id to be used in it
  <wh,dist,cust> = gen_order_id(w_id, d_id, c_id, order);

  total = 0;
  for(ord_item in order.items) {
    //Get the amount for the number of items ordered
    amount = get_amount(ord_item);
    total += amount;

    //Get the district information of the stock from the
    //supplier warehouse
    stock_info = get_dist_info_stock(ord_item);

    //Update the stock of the supplier warehouse
    update_stock(ord_item, amount);

    add("order_line", dist.order_id, w_id, d_id, stock_info,
        amount, ...);
  }
  total_pay = (1 + wh.tax + dist.tax) * total * (1 - cust.
    discount);
  return total_pay;
}

```

Figure 1: New\_order transaction

**Modeling hard to partition applications.** The model is ideally suited for applications where partitioning is implicit in the application logic. Fortunately, this property holds for a large class of HICCUP applications. Although applications that are hard to partition, e.g., social networks would not be a compelling use case for our model, there is no reason for them to perform worse using our model compared to RDBMS and key-value stores. On the contrary, the model provides an analytical playground to application developers to experiment with different partition layouts. This allows them to adopt the appropriate model (relational or key-value stores) for their application depending on the performance needs and development costs. The goal of the model is not to infer a partitioning layout for non-partitioned transactions but to provide a framework to the application developer to specify transactions with respect to a partitioning layout.

**System Initialization.** The model deterministically assigns transactions to logical partitions. Since the model always starts with an initial empty database state, the user application must provide initialization transactions to create an initial database state. Since those transactions would be created using the proposed model, data distribution would always be deterministic and unambiguous. It is important to mention that command logging to provide durability can be interpreted as an automatic program to generate a given desirable initial database state. This highlights the focus of the approach to logically partition code and not data.

#### 3.2 TPC-C Example

The ease of partitioning application logic will determine the ease of applicability of this approach. HICCUP applications that have simple application logic are well suited for this approach as we demonstrate by explaining our approach using the *new\_order* transaction class from the TPC-C benchmark as an example. We will not go into the detailed explanation of the benchmark which is available in its

```

txn new_order_update_stock(order) {
  result = <>;
  for(ord_item in order.items) {
    amount = get_amount(ord_item);
    stock_info = get_dist_info_stock(ord_item);
    update_stock(ord_item, amount);
    append(result, <stock_info, amount>);
  }
  return result;
}

```

**Figure 2: Transaction to read and update stock information**

```

PARTITIONING FUNCTION map(w_id) {return w_id;}
new_order PARTITION MAPPER map;
new_order_update_stock PARTITION MAPPER map;

```

**Figure 3: Mapping transactions to logical partitions**

specification<sup>2</sup>.

We begin by expressing the *new\_order* application logic without any notion of partitioning. This is followed by the transformation of the program into the proposed model by partitioning the *new\_order* transaction by warehouses. Finally, we show how developers can reason about the performance of this transaction and rewrite it in order to make it more efficient.

### 3.2.1 New\_Order Application Logic

The pseudocode for the *new\_order* transaction is outlined in Figure 1. The pseudocode preserves the data dependencies in the transaction class, but abstracts away other details of the implementation including data types. In the transaction logic, the order is first sanity checked and an order id is generated. For each item in the order, the amount for the number of ordered items is computed. The stock of the supplier warehouse is then updated and an order line is inserted for each item along with the district information of the stock. Finally, the total amount to be paid by the customer is computed and returned by the transaction. It is important to note that the fields retrieved by the *get\_dist\_info\_stock* method from the stock relation, which are used to insert a row in the *order\_line* relation, are never updated by the *update\_stock* method. The price for each item is present in the item relation, which is never updated and never grows as more warehouses are added.

The TPC-C benchmark was designed to scale in the number of warehouses, i.e., with additional warehouses, there are more customers who request greater number of *new\_order* transactions. Warehouses form the unit of distribution and scale-up in the benchmark. Many HICCUP applications have similar properties where a distribution unit leads to growth in both data and client requests and importantly, application logic is mostly affine to the distribution unit. *Existing programming models fail to take advantage of this property.* Our approach, which provides a high-level, intuitive, partitioned programming model with ACID guarantees, is ideally suited for this class of applications.

### 3.2.2 Code Partitioning by warehouses

In the TPC-C benchmark, a warehouse is the unit of logical partitioning. This follows intuitively because transaction

<sup>2</sup>[http://www.tpc.org/tpcc/spec/tpcc\\_current.pdf](http://www.tpc.org/tpcc/spec/tpcc_current.pdf)

```

txn new_order (w_id, d_id, c_id, order) {
  <wh,dist,cust> = gen_order_id(w_id, d_id, c_id, order);
  results [order.num_s_w_id]; i=0;

  for(s_id in order.supplier_w_id) {
    //Execute the remote transaction for all the order
    // items requested of the remote supplier warehouse
    results [i++] = EXEC new_order_update_stock
      (subset(order, s_id)) ON PARTITION (s_id);
  }

  //Use results to add order line and compute total pay
  total = 0;
  for(result in results) {
    for(item_result in result) {
      total += item_result.amount;
      add("order_line", dist.order_id, w_id, d_id,
        item_result, ...);
    }
  }
  total_pay = (1+wh.tax+dist.tax)*total*(1-cust.discount);
  return total_pay;
}

```

**Figure 4: New\_order transaction partitioned by warehouses**

logic is warehouse-affine and both data and number of client transactions increase with additional warehouses. In order to allow orders to be entirely supplied by a customer warehouse, each warehouse replicates the item relation.

Under this logical partitioning model, the *new\_order* transaction code is mostly local to the customer warehouse (*w\_id*). In order to execute the *get\_dist\_info\_stock* and *update\_stock* methods, the transaction needs to access a remote partition, since these methods operate on the supplier warehouse. We need to group these methods under a *new\_order\_update\_stock* transaction as shown in Figure 2, so that they can be invoked on the remote supplier warehouse logical partition.

In the partitioned model, each logical partition is identified by a unique logical partition identifier. Without loss of generality, the logical partitions can be identified by unique natural numbers. When a transaction is invoked, the logical partition where it would be executed must be specified, which establishes the relationship between the application logic unit and a logical partition in the model. This is done by using a mapping function that operates on the call parameters or any available context and returns an appropriate partition identifier. Each transaction class declares the mapping function it uses and defines it; however, all mapping functions output partition numbers in the same domain of logical partitions. For our example, warehouse id can be used as the logical partition identifier. The mapping function is declared and defined as shown in Figure 3. Now, a *new\_order* transaction can be invoked as:

**EXEC new\_order (w\_id, ...) ON PARTITION (w\_id)**

Finally, we re-partition the original *new\_order* transaction as shown in Figure 4, so that it is aware of logical partitions by invoking *new\_order\_update\_stock* on remote supplier warehouse logical partitions. Since transactions guarantee full ACID semantics even across logical partitions, the programmers do not have to worry about consistency issues.

### 3.2.3 Further Optimizations

A closer look at the partitioned *new\_order* transaction shown in Figure 4 exposes several performance issues ow-

```

txn new_order (w_id, d_id, c_id, order) {
  <wh,dist,cust> = gen_order_id(w_id, d_id, c_id, order);

  //Execute the remote transaction for all the order items
  //requested of the supplier warehouse in parallel
  PARALLEL EXEC new_order_update_stock
  (subset(order, s_id)) ON PARTITION (s_id)
  for s_id in order.supplier_w_id;

  total = 0;
  for(ord_item in order.items) {
    //Compute the pay amount
    amount = get_amount(ord_item);
    total += amount;
    //Get the district information of the stock using the
    //replicated stock table
    stock_info = get_dist_info_stock(ord_item);
    add("order_line", dist.order_id, w_id, d_id, stock_info,
        amount, ...);
  }
  total_pay = (1+wh.tax+dist.tax)*total*(1-cust.discount);
  return total_pay;
}

```

**Figure 5: New\_order transaction with further optimizations**

ing to control flow dependencies upon remote transaction results. There is a dependency on the results from the remote *new\_order\_update\_stock* transaction in order to insert order line entries. Since the item relation is already replicated on each logical partition, the amount can be computed locally. The real dependency of order line entries is to get the district information of the stock from the remote supplier warehouse. The district information fields are never updated so these fields can be replicated on each logical partition. Then, the *new\_order\_update\_stock* transaction can be modified so that it does not have to return a result and the *new\_order* transaction can be rewritten as shown in Figure 5. In order to invoke subtransactions in parallel, we provide **PARALLEL EXEC** construct. In order to use the construct, the transaction to be executed and the iterator to be used must be specified. Using this construct in Figure 5, “EXEC *new\_order\_update\_stock* (order) ON PARTITION (s\_id)” is invoked in parallel for all iterator elements (s\_id) in order.supplier\_w\_id list.

Since *new\_order* transactions constitute ~43% of the workload mix, this optimization might be worthwhile. As the number of warehouses increase, however, the size of the replicated stock relation with the district information fields increases as well. For a large number of warehouses, this optimization option may thus not be viable. Nevertheless, it is important to note that the model allowed the application developer to reason about the performance of the developed programs (*new\_order* programs) and improve them using the mentioned optimizations. Consequently, OLTP systems exposing a logical partitioning model can provide performance guarantees over varying classes of applications.

### 3.2.4 Summary

In a cluster-based setup, it is important to separate the concerns of writing efficient programs and ensuring their efficient execution by mapping to physical machines. This is done by exposing a programming model based on logical partitioning, which allows the application developers to reason about program performance. HICCUP applications,

which have an implicit distribution unit, can be intuitively modeled in the proposed programming model as shown in the TPC-C example. Strong global consistency semantics ensure ease of program development.

## 4. CHALLENGES AND ROADMAP

In this section, we outline the challenges that need to be tackled during the course of the project. We also provide some initial thoughts on how to address them.

### 4.1 Physical Implementation and Evaluation

**Logical to Physical Mapping.** In order to deploy a set of logical partitions over a set of physical partitions (physical machines with varying hardware characteristics), the logical partitions need to be mapped to the set of physical partitions. Each physical partition needs to run a partition executor that is responsible for handling transactions. This gives rise to the following sub-challenges, which we plan to investigate: (1) How to reuse a transaction processing system which meets performance demands on modern main-memory multi-core machines and extend it with the logical programming abstraction? (2) How to map a set of logical partitions to a set of physical partitions to ensure optimal resource usage of the cluster?

In order to meet the first challenge, we plan to use a state-of-the-art transaction processing system designed for main-memory multi-core machines [24], and extend it to operate in a distributed setup to provide the logical partitioning abstraction. In order to meet the second challenge, we plan to investigate cost models that attempt to approximate transaction execution costs in a physical cluster setup. This involves identifying and characterizing metrics in the cost model. The cost model can then be used to produce an initial mapping, which can be changed if workload characteristics change. We also plan to investigate transaction scheduling strategies without violating performance constraints, which would help in better resource utilization. **Local Concurrency Control and Global Commit.** One of the simplest models to run multiple transactions is to run them serially, which is what H-Store advocates [14]. This simple model will not work for distributed transactions unless a global schedule is pre-decided [23]. Running transactions serially would be a waste of compute power in multi-core architectures of today. There is also a need to hide memory latencies, network latencies and commit latencies. However, multi-threading with pessimistic concurrency control leads to contention bottlenecks in the lock manager [13] and multi-phase commits, which hurt performance.

In a partitioned model, distributed optimistic concurrency control holds promise owing to its parallel nature, lack of long critical sections and multi-phase commits [12]. It provides a global commit order in the distributed setup using local commit ordering decisions without necessitating a multi-phase protocol and hindering local transactions. We want to investigate the performance of a distributed optimistic concurrency control mechanism in order to support the transactional semantics of our approach. We also want to investigate how transaction aborts hurt performance under contention, and whether hybrid concurrency control and/or program repartitioning would improve system performance. **Evaluation.** In order to evaluate the system implementation, we plan to measure the scalability in transaction throughput, latency, abort rates and cost per transaction

with growing data and number of clients. We are particularly interested in evaluating these metrics for various physical machine configurations and logical-to-physical partition mappings. We plan to evaluate the system using the TPC-C benchmark. We are also interested in looking at other non-standard application benchmarks, which could provide clues on usability of the proposed model. We plan to compare the results with that of H-Store, a commercial RDBMS and Silo to characterize the system performance and understand the architecture design trade-offs.

## 4.2 Cloud Integration

The cloud computing infrastructure provides a flexible ecosystem for deployment of the system. In order to target the cloud computing infrastructure and satisfy varying HICCUP application requirements, we plan to investigate the necessary set of self-tuning tools to make the system cloud-ready. Deployment advisors that self-manage diverse, changing resources [26], adapt to the cloud infrastructure [27], have received a lot of attention lately. We plan to develop a deployment advisor that generates an optimal cluster configuration in the cloud from an application’s requirements (e.g., throughput, latency, cost constraints and goals) and programs by constructing and solving an optimization problem. We plan to investigate techniques to allow the deployment advisor to adapt to the variance and workload mix in the cloud infrastructure.

## 5. CONCLUSION

Designing main-memory OLTP systems that provide high performance, scalability and ease of development over a variety of HICCUP applications is an open challenge. In this paper, we identify the need to allow application developers to reason about the performance of their programs while maintaining ease of use. We propose an extension to the transactional abstraction to provide a high-level, intuitive, distributed logical programming model with strong global consistency semantics to the application developer. We have also shown how the model can be intuitively used by the application developer by using TPC-C as an example. We plan to evaluate the potential of this approach to meet the requirements of a variety of HICCUP applications.

## 6. REFERENCES

- [1] M. Armbrust, et al. A view of cloud computing. *Commun. ACM*, 53(4):50–58, 2010.
- [2] J. Baker, et al. Megastore: Providing scalable, highly available storage for interactive services. In *CIDR*, pages 223–234, 2011.
- [3] F. Chang, et al. Bigtable: A distributed storage system for structured data. *ACM Trans. Comput. Syst.*, 26(2), 2008.
- [4] A. Cheung, et al. Automatic partitioning of database applications. *PVLDB*, 5(11):1471–1482, 2012.
- [5] B. F. Cooper, et al. Pnuts: Yahoo!’s hosted data serving platform. *PVLDB*, 1(2):1277–1288, 2008.
- [6] J. C. Corbett, et al. Spanner: Google’s globally-distributed database. In *OSDI*, pages 261–264. USENIX Association, 2012.
- [7] C. Curino, et al. Schism: a workload-driven approach to database replication and partitioning. *PVLDB*, 3(1):48–57, 2010.
- [8] G. DeCandia, et al. Dynamo: amazon’s highly available key-value store. In *SOSP*, pages 205–220. ACM, 2007.
- [9] R. Escriva, B. Wong, and E. G. Sirer. Warp: Lightweight multi-key transactions for key-value stores. Technical report, Cornell University, Ithaca, 2013.
- [10] A. Floratou, et al. Can the elephants handle the nosql onslaught? *PVLDB*, 5(12):1712–1723, 2012.
- [11] P. Helland. Life beyond distributed transactions: an apostate’s opinion. In *CIDR*, pages 132–141, 2007.
- [12] M. Herlihy. Apologizing versus asking permission: Optimistic concurrency control for abstract data types. *ACM Trans. Database Syst.*, 15(1):96–124, 1990.
- [13] R. Johnson, et al. Shore-mt: a scalable storage manager for the multicore era. In *EDBT*, volume 360, pages 24–35. ACM, 2009.
- [14] R. Kallman, et al. H-store: a high-performance, distributed main memory transaction processing system. *PVLDB*, 1(2):1496–1499, 2008.
- [15] A. Lakshman and P. Malik. Cassandra: a decentralized structured storage system. *Operating Systems Review*, 44(2):35–40, 2010.
- [16] I. Pandis, et al. Data-oriented transaction execution. *PVLDB*, 3(1):928–939, 2010.
- [17] A. Pavlo, C. Curino, and S. B. Zdonik. Skew-aware automatic database partitioning in shared-nothing, parallel oltp systems. In *SIGMOD Conference*, pages 61–72, 2012.
- [18] D. Peng and F. Dabek. Large-scale incremental processing using distributed transactions and notifications. In *OSDI*, pages 251–264. USENIX Association, 2010.
- [19] J. Shute, et al. F1: A distributed sql database that scales. *PVLDB*, 6(11):1068–1079, 2013.
- [20] M. Stonebraker. New opportunities for new sql. *Commun. ACM*, 55(11):10–11, 2012.
- [21] M. Stonebraker, et al. The end of an architectural era (it’s time for a complete rewrite). In *VLDB*, pages 1150–1160. ACM, 2007.
- [22] D. Terry. Replicated data consistency explained through baseball. *Commun. ACM*, 56(12):82–89, 2013.
- [23] A. Thomson, et al. Calvin: fast distributed transactions for partitioned database systems. In *SIGMOD Conference*, pages 1–12, 2012.
- [24] S. Tu, et al. Speedy transactions in multicore in-memory databases. In *SOSP*, pages 18–32. ACM, 2013.
- [25] G. Weikum and G. Vossen. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery*. Morgan Kaufmann, San Francisco, CA, USA, 2002.
- [26] Q. Yin, et al. Rhizoma: A runtime for self-deploying, self-managing overlays. In *Middleware*, volume 5896, pages 184–204. Springer, 2009.
- [27] T. Zou, et al. Cloudia: A deployment advisor for public clouds. *PVLDB*, 6(2):109–120, 2012.