

Reactors: A Case for Predictable, Virtualized Actor Database Systems

Vivek Shah

University of Copenhagen, Denmark
bonii@di.ku.dk

Marcos Antonio Vaz Salles

University of Copenhagen, Denmark
vmarcos@di.ku.dk

ABSTRACT

The requirements for OLTP database systems are becoming ever more demanding. Domains such as finance and computer games increasingly mandate that developers be able to encode complex application logic and control transaction latencies in in-memory databases. At the same time, infrastructure engineers in these domains need to experiment with and deploy OLTP database architectures that ensure application scalability and maximize resource utilization in modern machines. In this paper, we propose a relational actor programming model for in-memory databases as a novel, holistic approach towards fulfilling these challenging requirements. Conceptually, relational actors, or *reactors* for short, are application-defined, isolated logical actors that encapsulate relations and process function calls asynchronously. Reactors ease reasoning about correctness by guaranteeing serializability of application-level function calls. In contrast to classic transactional models, however, reactors allow developers to take advantage of intra-transaction parallelism and state encapsulation in their applications to reduce latency and improve locality. Moreover, reactors enable a new degree of flexibility in database deployment. We present REACTDB, a system design exposing reactors that allows for flexible virtualization of database architecture between the extremes of shared-nothing and shared-everything without changes to application code. Our experiments illustrate latency control, low overhead, and asynchronicity trade-offs with REACTDB in OLTP benchmarks.

ACM Reference Format:

Vivek Shah and Marcos Antonio Vaz Salles. 2018. Reactors: A Case for Predictable, Virtualized Actor Database Systems. In *SIGMOD'18: 2018 International Conference on Management of Data, June 10–15, 2018, Houston, TX, USA*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3183713.3183752>

1 INTRODUCTION

Three trends are transforming the landscape of OLTP systems. First, a host of latency-sensitive OLTP applications has emerged in areas as diverse as computer games, high-performance trading, and the web [12, 49, 55]. This trend brings about challenging performance requirements, including mechanisms to allow developers to reason about transaction latencies and scalability of their applications with large data and request volumes [48, 53]. Second, database

systems are moving towards solid state, in particular in-memory storage [32], and hardware systems are integrating increasingly more cores in a single machine. This trend brings about new requirements for database architecture, such as processing efficiency in multi-core machines and careful design of concurrency control strategies [52, 57]. Third, there is a need to operate databases out of virtualized infrastructures with high resource efficiency [8, 31, 40]. This trend leads to the requirement that virtualization abstractions for databases impose low overhead and allow for flexible deployments without causing changes to application programs.

Recent research in OLTP databases has shown that addressing all of these requirements is a hard problem. On the one hand, shared-nothing database designs, such as those of H-Store [50] or HyPer [30], fail to provide appropriately for multi-core efficiency in the presence of cross-partition transactions [20, 52]. This is due to the impact of overheads in mapping partitions to cores and of synchronous communication in distributed transactions across partitions. Consequently, transaction throughput and latencies in these systems are very sensitive to how data is partitioned. On the other hand, shared-everything databases have a hard time achieving multi-core scalability. To do so, these systems either internally partition their data structures, e.g., DORA [42], or benefit from affinity of memory accesses to cores in transactions, e.g., Silo [52]. Thus, deployment decisions can affect efficiency and scalability in these systems and are difficult to get right across application classes.

As a consequence, both developers and infrastructure engineers in demanding OLTP domains have a hard time controlling the performance of their transactional databases. Despite advances in profiling tools to identify causes of latency variance in database systems [29], today developers lack clear abstractions to reason at a high level about the interplay of complex, potentially parallelizable application logic and transaction latencies. In addition, the variety of modern in-memory database engines, including numerous specialized designs ranging internally from shared-nothing to shared-everything [43, 47, 56], challenges the ability of infrastructure engineers to flexibly experiment with and adapt database architecture without affecting application code.

Actor programming models provide desirable primitives for concurrent and distributed programming [2, 4, 27], which of late have evoked a strong interest in the database community [9]. To holistically meet the challenging requirements imposed on OLTP systems, we propose a new actor programming model in relational databases called *Relational Actors* (or *reactors* for short). Reactors are special types of actors that model logical computational entities encapsulating state abstracted as relations. For example, reactors can represent application-level scaling units such as accounts in a banking application or warehouses in a retail management application. Within a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMOD'18, June 10–15, 2018, Houston, TX, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-4703-7/18/06...\$15.00

<https://doi.org/10.1145/3183713.3183752>

reactor, developers can take advantage of classic database programming features such as declarative querying over the encapsulated relations. To operate on state logically distributed across reactors, however, developers employ explicit asynchronous function calls. The latter allows developers of latency-sensitive OLTP applications to write their programs so as to minimize cross-reactor accesses or overlap communication with computation. Still, a transaction across multiple reactors provides serializability guarantees as in traditional databases, thus relieving developers from struggling with complex concurrency issues. Reactors allow application-level modeling between the extremes of a relational database (single reactor encapsulating all relations) and key-value stores (each reactor encapsulating a key-value pair).

To address the challenges of architectural flexibility and high resource efficiency in multi-core machines, we design an in-memory database system that exposes reactors as a programming model. This system, called REACTDB (RElational ACTor DataBase), decomposes and virtualizes the notions of sharing of compute and memory in database architecture. First, we introduce a database containerization scheme to enclose shared-memory regions in a machine, each storing state for one or many reactors. Second, within a container, compute resources abstracted as transaction executors can be deployed to either share or own reactors. The combination of these two notions allows infrastructure engineers to experiment with deployments capturing a range of database architecture patterns by simply changing a configuration file. At the same time, no changes are required to application code using reactors.

Example: Digital Currency Exchange. We abstract an application coded using a set of reactors as a *reactor database*. Consider a simplified digital currency exchange application, in which users may buy or sell currency through their credit card providers. Figure 1 contrasts how such an application would be written with a classic transactional database and a reactor database in parts (a) and (b), respectively. The exchange wishes to limit its settlement risk from cancelled credit card payments. To do so, it follows two application rules: (1) stop accepting orders if any single provider’s total unsettled exposure goes above the `p_exposure` threshold in relation `settlement_risk`; (2) reject orders that cause the total risk-adjusted exposure across all providers to exceed the `g_risk` threshold. In part (a), this logic is encoded in a stored procedure that computes risk-adjusted exposure through an expensive function `sim_risk` and caches its result for a time period. In part (b), the same logic is expressed with reactors. The exchange and each of the providers are modeled as relational actors with private state (relations) that can execute certain procedures. The Exchange reactor can execute `auth_pay` and encapsulates information about provider names and the settlement limits. Provider reactors store information from risk-adjusted exposure calculations per provider as well as fragments of the relation `orders` with the payments for each provider, and can execute procedures `calc_risk` and `add_entry`. In `auth_pay`, the Exchange reactor invokes asynchronous calls to Provider reactors, making explicit the available intra-transaction parallelism. Since the exchange strives for the lowest latency possible, this program formulation improves transaction response time with respect to part (a) in a way that is clearly explainable to a developer pursuing application performance optimization. In addition, it becomes explicit in transaction programs that code is conceptually

moved close to the data it touches, allowing developers to control for locality. At the same time, ACID properties are guaranteed for `auth_pay`, despite asynchronous calls to a function, `calc_risk`, that includes database updates, user-defined abort conditions, and potentially even nondeterminism in the calculation of `sim_risk` [5]. We further elaborate on this example and our model in Section 2.

How are reactors different from database partitioning?

In contrast to database partitioning, which is a data-oriented optimization, reactors represent a *compute-oriented abstraction*. In the example above, reactors can be used to model horizontal partitioning of tuples, vertical partitioning of relations, or any arbitrary grouping of relation fragments. In addition, reactors allow for functional decomposition and modeling of affinity and parallelism in arbitrary application logic. In the example, detecting and fully exploiting the parallelism of `auth_pay` in the stored procedure formulation of part (a) in a classic relational database would require complex control-flow analysis, and may not be possible at all [45].

Contributions. In short, we make the following contributions:

- (1) We present a novel logical abstraction for relational databases called *reactors*. This abstraction is grounded on transactional semantics offering serializability and an asynchronous programming model that allows for *encoding complex application logic while considering relative latency of different transactional programs* (Section 2).
- (2) We discuss the design of REACTDB, an in-memory database system exposing reactors. REACTDB enables *configuration of database architecture at deployment time* in a multi-core machine without changes to application code (Section 3).
- (3) In experiments with classic OLTP benchmarks, reactors provide latency control at the microsecond scale for varied program formulations. In addition, for given program formulations, database architecture can be configured to control the trade-off between asynchronicity and load (Section 4).

2 PROGRAMMING MODEL

2.1 Reactor Concepts

In contrast to classic transactional or actor models, reactors bring together all of the following concepts:

- (1) A reactor is an application-defined logical actor that encapsulates state abstracted using relations.
- (2) Declarative queries are supported only on a single reactor. Communication across reactors is achieved by asynchronous function calls. A computation (function) across reactors consists of a sequence of intra-reactor statements and/or nested cross-reactor function calls.
- (3) Computations across reactors have transactional guarantees.
- (4) Reactors provide an abstract computational cost model for reflecting on relative latency across program formulations.

2.2 Programming with Reactors

2.2.1 Application-Defined Relational Actors. A *reactor* is an actor specialized for the management of state abstracted by the relational model. The pseudocode in Figure 2 conceptualizes the capabilities of a reactor. As a regular actor [2], a reactor encapsulates a state, which can be accessed by computations invoked on the reactor. However, unlike in a regular actor, in which communication is

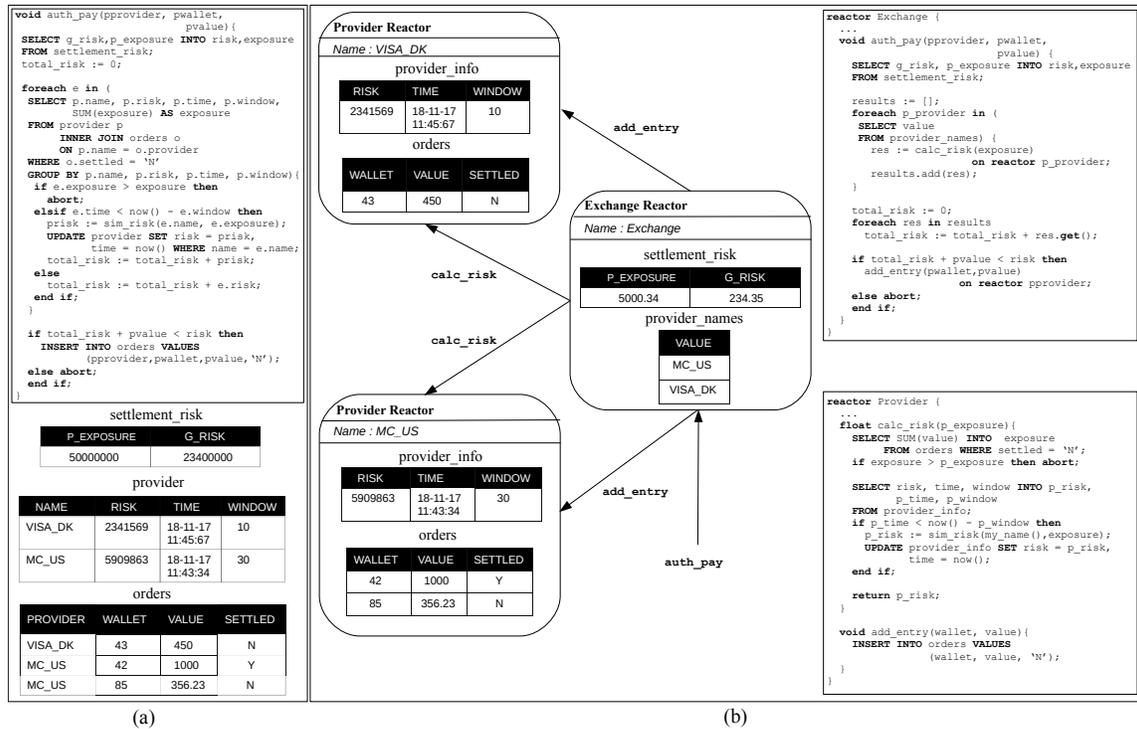


Figure 1: A simplified currency exchange application in: (a) the classic transactional model, and (b) the reactor model.

```

Reactor : Actor {
  RelationalState rs;

  Future execute(compute_fn, args) {
    return new Future(compute_fn(args, rs));
  }
}

```

Figure 2: Conceptual view of reactors as an actor type.

typically achieved by non-blocking send and blocking receive primitives, the only form of communication with a reactor is through *asynchronous function calls* returning promises [37]. Moreover, the code of such functions is similar to that of database stored procedures, which can intermix declarative queries over relations with other program logic and function calls. These asynchronous function calls are abstracted in Figure 2 by the `execute` function, which takes as an argument a function to be computed on the reactor’s state along with appropriate arguments, and returns a promise representing the result of the computation. In the remainder of this paper, we refer to such a result as a *future*, and use the terms *function* and *procedure* on a reactor interchangeably.

A *reactor database* is a collection of reactors, each of which must abide by a reactor type. Reactor types, defined by application developers, specify the functions that should be invoked in a reactor and determine the relation schemas encapsulated in the reactor state. To instantiate a reactor database, we need to declare the names and types of the reactors constituting the database and a schema creation function for each reactor type. The developer cannot create or destroy reactors; these purely logical entities are accessible by their declared names for the lifetime of the application, bound by the failure model of the reactor database. In contrast to objects in an object-oriented database [13, 34], reactors are active computational

entities, i.e., a reactor is a combination of a logical thread of control and an encapsulated relational state accessible *exclusively* by that logical thread. While objects encapsulate complex types, reactors encapsulate whole relational schemas; declarative querying happens only within, not across reactors, and communication across reactors is explicit through asynchronous function calls.

In the example of Figure 1(b), the state of a Provider reactor results from both horizontal and vertical fragmentation of the original providers relation from part (a) as well as horizontal fragmentation of the orders relation (with omission of the provider column). The state of the Exchange reactor retains a relation `provider_names`, since the latter is necessary for accessing Provider reactors, and the `settlement_risk` relation. This illustrates that different reactors may contain either the same or different schemas, according to their types.

It is not necessary to know in advance all the providers and their names to model the reactor database. It is sufficient to know: (1) the types of the reactors expected, namely Exchange and Provider; (2) the schemas and functions of each reactor type; (3) the name mapping to address provider reactors. As such, adding new providers does not necessitate rewriting the application logic.

2.2.2 Asynchronous Function Calls. To invoke a procedure on a reactor, we must explicitly use the declared name of the reactor where the computation must be executed. The procedure logic can access the relational state on the reactor where it is invoked through declarative queries. If the procedure needs to access the state of another reactor, then it must invoke another procedure on the target reactor. This is necessary because the states of different reactors are disjoint. Since the result of a procedure is represented

by a future, the calling code can choose to wait for the result of the future, invoke procedures on other reactors, or execute further application logic. This flexibility allows application developers to expose parallelism within the procedure.

In Figure 1(b), the syntax `procedure_name(args) on reactor reactor_name` specifies an asynchronous procedure call routed to the reactor with a given name. In the logic of `auth_pay`, the execution of `calc_risk` is overlapped on each of the Provider reactors and the futures returned by the calls are stored in the `results` list. The Exchange reactor sums up the total risk by accessing the value of each future by invoking `get()` on the future object. If the total risk is within the allowed risk, then the exchange reactor performs another nested asynchronous procedure call to `add_entry` on the provider reactor with name given as a parameter to `auth_pay`. This call results in adding an order at the target provider reactor.

Does the reactor programming model necessitate manual optimization? We posit that reactors can act as a bridging model between a classic database abstraction and a key-value API. Reactors provide the possibility for developers to navigate the extremes between a single-reactor database with full SQL support and a radical decomposition of individual tuples as reactors. We envision that the typical application modeling will be a hybrid between these two extremes balancing reuse of query optimization machinery with low-level performance control. The popularity of NoSQL databases points to the need and willingness among application developers to obtain higher performance control and scalability for their applications even at the cost of sacrificing traditional database features such as query optimization and transaction support.

2.2.3 Reactor Consistency using Transactional Semantics. To guarantee consistency of the state encapsulated by a reactor database, the semantics of procedure invocations on reactors is transactional. We differentiate between top-level and nested asynchronous procedure calls. Top-level calls are executed by clients on a reactor and are termed interchangeably *transactions* or *root transactions*. Transactions respect the classic ACID properties [10]. We denote a concrete execution i of a transaction by T_i .

Nested asynchronous procedure calls are executed by a reactor on another reactor. Since these calls must always occur within the overall context of a root transaction, they are called *sub-transactions*. We denote a concrete execution of a sub-transaction j of transaction T_i on a reactor k by $ST_{i,j}^k$. Sub-transactions allow programmers to structure their computations for performance, allowing for concurrent computation on (logically) distributed state among reactors. Sub-transactions are *not* used, however, to allow for partial commitment. Any condition leading to an abort in a sub-transaction leads to the abort of the corresponding root transaction. This approach towards the semantics of nested calls is exactly the reverse of what is adopted in classic systems such as Argus [36], reflecting our focus on leveraging the high degree of *physical parallelism* in modern commodity hardware for transactional processing as opposed to managing faults in settings with a high degree of physical distribution (e.g., geo-distribution) as in previous work. A transaction or sub-transaction completes only when all its nested sub-transactions complete. This frees the client logic from explicitly synchronizing on the result of a sub-transaction invocation if it does not need the result of the sub-transaction.

For example in Figure 1(b), the `auth_pay` procedure does not wait on the result of the `add_entry` procedure call, since the programming model guarantees that the transaction corresponding to `auth_pay` only completes when all its sub-transactions complete. In addition, `auth_pay` aborts if any of the asynchronous calls to `calc_risk` raises an abort due to excessive provider exposure.

Any program in the classic transactional model can be trivially remodeled in the reactor programming model by specifying a single reactor. For example, we could model a single exchange reactor with the schema and application logic shown in Figure 1(a). However, the benefits of our programming model are only fully achieved when developers of latency-sensitive OLTP applications remodel their logic as done in Figure 1(b). In particular, in the reformulated logic, intra-transaction parallelism is exposed. Furthermore, the trade-off between scalability on the number of provider reactors and latency of executing the logic of `auth_pay` becomes explicit.

Is there a development methodology to architect an application using reactors? An interesting avenue for future research is to explore an analytical machinery for modeling and comparing the quality of reactor database designs, similar to the entity relationship model and decomposition of universal relations by functional dependencies in classic relational databases. Although answering this question is beyond the scope of this paper, we envision that developers could start from a single reactor type with the whole relational schema and all application functions. Through an iterative process of performance analysis and decomposition of application functionality into multiple reactor types, developers can improve latency of their applications through cross-reactor communication, and also identify inherent scalability limitations by analyzing the degree of locality in application logic.

2.2.4 Intra-Transaction Safety. Introducing asynchrony in a transactional abstraction is not trivial. Since asynchronicity exposes intra-transaction parallelism, race conditions could arise when sub-transactions that conflict on a data item are invoked asynchronously on the same reactor. Moreover, such invocations would violate the illusion that a reactor is a computational entity with a single logical thread of control. To avoid these issues, we must enforce that at most one execution context is active for a given reactor and root transaction at any time.

First, we enforce that whenever a reactor running a procedure directly executes a nested procedure invocation on itself, the nested invocation is executed *synchronously*. This policy corresponds to inlining the sub-transaction call, resulting in future results being immediately available. To deal with nested asynchronous invocations, we define the *active set* of a reactor k as the set of sub-transactions, regardless of corresponding root transaction, that are currently being executed on reactor k , i.e., have been invoked, but have not completed. Thus, the runtime system must conservatively disallow execution of a sub-transaction $ST_{i,j}^k$ when:

$$\exists ST_{i,j'}^k \in \text{active_set}(k) \wedge j' \neq j$$

This dynamic safety condition prohibits programs with cyclic execution structures across reactors and programs in which different paths of asynchronous calls lead to concurrent sub-transactions on the same reactor. By introducing this safety condition, the runtime conservatively assumes that conflicts may arise in asynchronous

accesses to the same reactor state within a transaction, and thus aborts any transaction with such dangerous structures.

By leveraging this dynamic safety condition, we envision that appropriate testing of transaction logic at development time will be sufficient to root out most, if not all, dangerous structures from the code of latency-sensitive OLTP applications. However, formalizing static program checks to aid in detection of dangerous call structures among reactors is an interesting direction for future work.

2.3 Conflict-Serializability of Transactions

To formalize the correctness of concurrent executions of transactions in reactors, we show equivalence of serializable histories in the reactor model to serializable histories in the classic transactional model. We restrict ourselves exclusively to the notion of conflict-serializability. Technically, our formalization is similar to reasoning on nonlayered object transaction models [54].

2.3.1 Background. We first review the formalism introduced by Bernstein et al. [10, page 27] for the classic transactional model and introduce relevant notation. In this model, the database consists of a collection of named data items, and transactions encapsulate a sequence of operations. A transaction T_i is formalized as a partial ordering of operations with an ordering relation $<_i$ and comprises a set of operations. Operations include reads and writes, along with either a commit or an abort. A read from a data item x is denoted $r_i[x]$, a write to x denoted $w_i[x]$, while a commit is denoted c_i and an abort a_i . The ordering relation $<_i$ orders *conflicts*. Two operations conflict iff at least one of them is a write and both of them reference the same named item. We assume that a transaction does not contain multiple operations of the same type to the same named data item as in [10, page 27] without any impact on the results.

2.3.2 Reactor Model. Without loss of generality, we assume reactor names and sub-transaction identifiers to be drawn from the set of natural numbers. Recall that we denote a sub-transaction j in transaction T_i on reactor k by $ST_{i,j}^k$, $r_{i,j}^k[x]$ denotes a read from data item x , and $w_{i,j}^k[x]$ denotes a write to data item x in $ST_{i,j}^k$. Note that data items in different reactors are disjoint. Using this notation, we can define a sub-transaction in the *reactor model* as follows.

Definition 2.1. A sub-transaction $ST_{i,j}^k$ is a partial order with ordering relation $<_{i,j}$ where,

- (1) $ST_{i,j}^k \subseteq \{ r_{i,j}^k[x], w_{i,j}^k[x], ST_{i,j'}^{k'} \mid x \text{ is a data item in } k, j' \text{ is a sub-transaction identifier s.t. } j' \neq j \}$;

- (2) Let

$$\text{basic_ops}(r_{i,j}^k[x]) = \{ r_{i,j}^k[x] \}$$

$$\text{basic_ops}(w_{i,j}^k[x]) = \{ w_{i,j}^k[x] \}$$

$$\text{basic_ops}(ST_{i,j}^k) = \{ \text{basic_ops}(o) \mid o \in ST_{i,j}^k \},$$

$$\text{if } o_1 \in ST_{i,j}^k \wedge o_2 \in ST_{i,j}^k \\ \wedge r_{i,j'}^{k'}[x] \in \text{basic_ops}(o_1)$$

$$\wedge w_{i,j''}^{k''}[x] \in \text{basic_ops}(o_2),$$

then either $o_1 <_{i,j} o_2$ or $o_2 <_{i,j} o_1$.

Note that the ordering relation $<_{i,j}$ of a sub-transaction establishes order according to conflicts in leaf-level basic operations,

potentially nested in sub-transactions. These leaf-level basic operations are determined by the function `basic_ops` in the definition.

Definition 2.2. A transaction T_i is a partial order with ordering relation $<_i$ where,

- (1) $T_i \subseteq \{ ST_{i,j}^k \} \cup \{ a_i, c_i \}$;

- (2) $a_i \in T_i \iff c_i \notin T_i$;

- (3) if t is c_i or a_i (whichever is in T_i), then for every other operation $p \in T_i, p <_i t$;

- (4) if $o_1 \in T_i \wedge o_2 \in T_i \wedge o_1, o_2 \notin \{ a_i, c_i \} \wedge r_{i,j'}^{k'}[x] \in \text{basic_ops}(o_1) \wedge w_{i,j''}^{k''}[x] \in \text{basic_ops}(o_2)$, then either $o_1 <_i o_2$ or $o_2 <_i o_1$.

Formally, a transaction comprises exclusively sub-transactions, and the relation $<_i$ orders sub-transactions according to conflicts in their nested basic operations. In the reactor model, two sub-transactions *conflict* iff the basic operations of at least one of them contain a write and the basic operations of both of them reference the same named item in the same reactor. Under this extended notion of a conflict, the definition of history, serial history, equivalence of histories and serializable history in the reactor model are the same as their definitions in the classic transactional model [10], but with sub-transactions replacing basic operations. Similar to nested transaction models [6], we then wish to establish an equivalence between serializability of transactions in the reactor model and serializability in the classic transactional model. To do so, we proceed by defining an appropriate projection of the reactor model into the classic transactional model.

Definition 2.3. The projection of a basic operation o from the reactor model to the classic transactional model, denoted by $P(o)$, is defined as:

- (1) $P(r_{i,j}^k[x]) = r_i[k \circ x]$

- (2) $P(w_{i,j}^k[x]) = w_i[k \circ x]$

- (3) $P(c_i) = c_i$

- (4) $P(a_i) = a_i$

where \circ denotes concatenation.

The definition provides a name mapping from the disjoint address spaces of reactors to a single address space, which is done by concatenating the reactor identifier with name for a data item.

Definition 2.4. The projection of a sub-transaction $ST_{i,j}^k$ from the reactor model to the classic transactional model, denoted by $P_S(ST_{i,j}^k)$ is a partial order with ordering relation $<_S^{i,j}$:

- (1) $P_S(ST_{i,j}^k) \subseteq \{ P(o) \mid o \in \text{basic_ops}(ST_{i,j}^k) \}$;

- (2) if $o_1 \in ST_{i,j}^k \wedge o_2 \in ST_{i,j}^k \wedge o_1 <_{i,j} o_2 \wedge o_1, o_2$ are reads or writes, then $P(o_1) <_S^{i,j} P(o_2)$;

- (3) if $ST_{i,j'}^{k'} \in ST_{i,j}^k$, then $<_S^{i,j}$ is extended by $<_S^{i,j'}$;

- (4) if $o_1 \in ST_{i,j}^k \wedge o_1$ is a read or a write $\wedge ST_{i,j'}^{k'} \in ST_{i,j}^k \wedge o_1 <_{i,j} ST_{i,j'}^{k'}$, then $P(o_1) <_S^{i,j} P_S(ST_{i,j'}^{k'})$;

- (5) if $ST_{i,j'}^{k'} \in ST_{i,j}^k \wedge o_2 \in ST_{i,j}^k \wedge o_2$ is a read or a write $\wedge ST_{i,j'}^{k'} <_{i,j} o_2$, then $P_S(ST_{i,j'}^{k'}) <_S^{i,j} P(o_2)$;

- (6) if $ST_{i,j'}^{k'} \in ST_{i,j}^k \wedge ST_{i,j''}^{k''} \in ST_{i,j}^k \wedge ST_{i,j'}^{k'} <_{i,j} ST_{i,j''}^{k''}$, then $P_S(ST_{i,j'}^{k'}) <_S^{i,j} P_S(ST_{i,j''}^{k''})$.

Definition 2.5. The projection of a transaction T_i from the reactor model to the classic transactional model, denoted by $P_T(T_i)$ is a partial order with ordering relation $<_T^i$:

- (1) $P_T(T_i) \subseteq (\bigcup_{ST_{i,j}^k \in T_i} P_S(ST_{i,j}^k)) \cup \{P(o) \mid o \in T_i \wedge o \text{ is a commit or abort}\}$;
- (2) $<_T^i$ is extended by $\bigcup_{ST_{i,j}^k \in T_i} <_S^{i,j}$;
- (3) if $ST_{i,j}^k \in T_i \wedge ST_{i,j'}^{k'} \in T_i \wedge ST_{i,j}^k <_i ST_{i,j'}^{k'} \wedge o_1 \in P_S(ST_{i,j}^k) \wedge o_2 \in P_S(ST_{i,j'}^{k'})$, then $o_1 <_T^i o_2$;
- (4) if t is c_i or a_i (whichever is in $P_T(T_i)$), for any further operation $p \in P_T(T_i)$, $p <_T^i t$.

The definitions unroll all sub-transactions in the reactor model into read and write operations in the classic transactional model while maintaining ordering constraints.

Definition 2.6. The projection $P(H)$ of a history H over a set of transactions $T = \{T_1, T_2, \dots, T_n\}$ from the reactor model to the classic transactional model is a partial order with ordering relation $<_{P_H}$ over a set of transactions $T' = \{P_T(T_1), P_T(T_2), \dots, P_T(T_n)\}$ iff:

- (1) $P_{P_H}(H) = \bigcup_{i=1}^n P_T(T_i)$;
- (2) $<_{P_H}$ is extended by $\bigcup_{i=1}^n <_T^i$;
- (3) if $o_1 \in P_S(ST_{i,j}^k) \wedge o_2 \in P_S(ST_{i',j'}^{k'}) \wedge ST_{i,j}^k \in T_i \wedge ST_{i',j'}^{k'} \in T_{i'} \wedge ST_{i,j}^k <_H ST_{i',j'}^{k'}$, then $o_1 <_{P_H} o_2$ as long as o_1 and o_2 conflict.

THEOREM 2.7. *A history H is serializable in the reactor model iff its projection $H' = P(H)$ in the classic transactional model is serializable.*

PROOF. See Appendix A. □

Theorem 2.7 implies that, with appropriate care, a scheduler for the classic transactional model can be used to implement one for the reactor model. In Section 3, we reuse an optimistic concurrency control (OCC) scheduler [52] and two-phase commit (2PC) [10].

2.4 Computational Cost Model

In this section, we introduce a cost model to support developers of latency-sensitive applications in controlling the latency of a transaction program expressed using reactors. Clearly, latency depends heavily on program structure. For example, certain programs can overlap asynchronous invocations of functions in other reactors with processing logic; other programs may do so only conditionally, or have data dependencies between different asynchronous function calls. For concreteness, we focus on a subset of programs modeled after a fork-join parallelism pattern [35].

Fork-Join Sub-Transactions. We call a sub-transaction fork-join if it comprises: (a) sequential logic, potentially involving synchronous calls to sub-transactions; (b) parallel logic consisting of calls to children fork-join sub-transactions such that all asynchronous invocations happen simultaneously at one given program point only, are potentially overlapped with synchronous logic, and then all future results are collected. Consider one such sub-transaction $ST_{i,j}^k$. We call $\text{sync}_{seq}(ST_{i,j}^k)$ its sequence of children sub-transactions and $P_{seq}(ST_{i,j}^k)$ its processing logic executed sequentially. To capture communication costs, we term $C_s(k, k')$ the cost to send a sub-transaction call from reactor k to reactor k' , and $C_r(k', k)$ the

$$\begin{aligned}
L(ST_{i,j}^k) = & P_{seq}(ST_{i,j}^k) + \sum_{ST_{i,j'}^{k'} \in \text{sync}_{seq}(ST_{i,j}^k)} L(ST_{i,j'}^{k'}) \\
& + \sum_{k' \in \text{dest}(\text{sync}_{seq}(ST_{i,j}^k))} (C_s(k, k') + C_r(k', k)) \\
& + \max \left(\max_{ST_{i,j'}^{k'} \in \text{async}(ST_{i,j}^k)} \left(L(ST_{i,j'}^{k'}) + C_r(k', k) \right) \right. \\
& \left. + \sum_{k'' \in \text{dest}(\text{prefix}(\text{async}(ST_{i,j}^k), ST_{i,j'}^{k'}))} C_s(k, k'') \right), \\
& P_{ovp}(ST_{i,j}^k) + \sum_{ST_{i,j'}^{k'} \in \text{sync}_{ovp}(ST_{i,j}^k)} L(ST_{i,j'}^{k'}) + \\
& \left. \sum_{k' \in \text{dest}(\text{sync}_{ovp}(ST_{i,j}^k))} (C_s(k, k') + C_r(k', k)) \right)
\end{aligned}$$

Figure 3: Modeling latency cost of a fork-join sub-transaction in the reactor model.

cost to receive a result from k' at k . The sequence of children sub-transactions of $ST_{i,j}^k$, executed asynchronously are denoted $\text{async}(ST_{i,j}^k)$. The synchronous children sub-transactions and processing logic overlapped with the asynchronous sub-transactions are represented by $\text{sync}_{ovp}(ST_{i,j}^k)$ and $P_{ovp}(ST_{i,j}^k)$, respectively. Given a sequence S of sub-transactions, $\text{prefix}(S, ST_{i,j}^k)$ denotes the sequence of sub-transactions in S up to $ST_{i,j}^k$ and including it. Moreover, we say that $\text{dest}(ST_1, \dots, ST_n)$ represents the sequence of reactors that sub-transactions ST_1, \dots, ST_n execute on.

Now, the latency cost of $ST_{i,j}^k$ is modeled by the formula in Figure 3, assuming the parallelism in asynchronous sub-transactions is fully realized. The same formula can be applied recursively to compute the latency cost for sub-transactions of arbitrary depth. Since a root transaction is a special case of a sub-transaction, i.e., a sub-transaction without a parent, the same formula applies, modulo any overheads incurred for commitment.

Uses. Many applications can be written in the reactor model with fork-join sub-transactions. Consider the transaction `auth_pay` in Figure 1(b). While it is not fork-join as presented, it can be easily decomposed into two sequentially invoked fork-join sub-transactions, namely one to calculate `total_risk` and one to conditionally call `add_entry`. All the benchmarks evaluated in our experiments can also be expressed with fork-join sub-transactions. Notwithstanding, the reactor programming model allows for any synchronization pattern with futures respecting the conditions of Section 2.2.4, and is *not limited* to fork-join sub-transactions. The cost model of Figure 3 can be extended to cover other program classes, if necessary.

We envision that developers may employ cost modeling as in Figure 3 in a way similar to algorithmic complexity measures, e.g., [1], to compare alternative formulations for transaction programs. Developers can improve the latency of their sub-transactions by: (1) increasing asynchronicity of children sub-transactions, (2) overlapping execution of application logic by introducing sub-transactions, and (3) reducing the processing cost of the application logic.

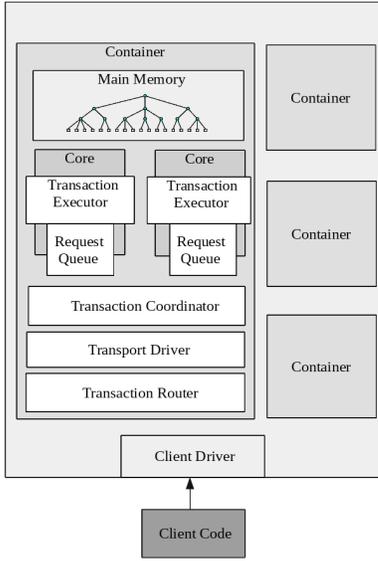


Figure 4: REACTDB’s architecture.

Finally, developers can reason about scalability by considering how data, computation and communication are spread among reactors. In particular, if developers architect their applications such that increasing amounts of data and computation are distributed among increasing numbers of reactors while at the same time keeping the number of cross-reactor calls roughly constant per transaction, then adequate transactional scalability should be expected.

Limitations. Again similar to algorithmic complexity measures, a cost model such as the one in Figure 3 is exposed to limitations in the estimation of its various parameters for concrete fork-join transaction programs and system realizations of the reactor model. It may be impossible to estimate reliably how costly processing will be or how many sub-transaction invocations will fall into $\text{sync}_{seq}(ST_{i,j}^k)$, $\text{sync}_{ovp}(ST_{i,j}^k)$, and $\text{async}(ST_{i,j}^k)$, as these parameters may be determined by complex program logic and data dependencies. In addition, the concrete system realization may not express the full parallelism encoded in the program. Moreover, determination of cost model parameters such as C_s and C_r may be compromised by measurement errors. Finally, the cost model considers transaction programs in isolation, and does not include interference or queueing effects. Notwithstanding, we remark in our experimental evaluation that under certain conditions, the cost model can closely capture latency differences between alternative transaction programs.

3 SYSTEM ARCHITECTURE

3.1 Overview

In this section, we discuss the architecture of REACTDB, an in-memory database system that exposes the reactor programming model. The design of REACTDB aims at providing control over the mapping of reactors to physical computational resources and memory regions under concurrency control. The system implementation currently targets a single multi-core machine for deployment; however, REACTDB’s architecture is designed to allow for deployments in a cluster of machines, which we leave for future work.

As shown in Figure 4, REACTDB’s architecture is organized as a collection of database *containers*. A container abstracts a (portion of a) machine with its own storage (main memory) and associated mechanisms for transactional consistency. Each container is isolated and does not share its data with other containers. Containers are associated with computational resources (cores) disjoint from other containers, abstracted by *transaction executors*. A transaction executor consists of a thread pool and a request queue, and is responsible for executing requests, namely asynchronous procedure calls. Each transaction executor is pinned to a core. Single-container transactions are managed by the concurrency control mechanism within the container, while a *transaction coordinator* runs a commitment protocol for transactions spanning multiple containers. Transactional durability is currently unavailable in our implementation, but could be achieved by a combination of techniques such as fast log-based recovery [59] and distributed checkpoints [24]. Alternatively, an approach such as FaRM’s could be employed to minimize any impact of durability mechanisms on latency [22].

Each container stores a two-level mapping between a reactor and a transaction executor. On the first level, a reactor is mapped to one and only one container. Together with appropriate container deployment, this constraint ensures that asymmetrically large communication costs are only introduced between, but not within, reactors, in line with our computational cost model. On the second level, a reactor can be mapped to one or more transaction executors in a container. *Transaction routers* decide the transaction executor that should run a transaction or sub-transaction according to a given policy, e.g., round-robin or affinity-based.

Transport drivers handle communication across containers. REACTDB has a driver component that is used by client code to send transactions into the system for processing. REACTDB accepts pre-compiled stored procedures written in the reactor programming model in C++ against a record manager interface. An instance of a pre-compiled stored procedure and its inputs forms a transaction.

3.2 Concurrency Control

3.2.1 Single Container Transactions. Every transaction or sub-transaction written in the reactor programming model specifies the reactor where it must be executed. If the destination reactor of a child sub-transaction is hosted in the same container as the parent sub-transaction, the child sub-transaction is executed synchronously within the same transaction executor to minimize the communication overhead of migrating across transaction executors. If all the sub-transactions in the execution context of a root transaction are executed within one container, then the native concurrency control mechanism of the container is used to guarantee serializability. As a consequence of Theorem 2.7, REACTDB can reuse an existing concurrency control mechanism, and we chose Silo’s high-performance OCC implementation [52].

3.2.2 Multi-Container Transactions. When a sub-transaction is invoked on a reactor mapped to a container different than the current container, the call is routed by the transport driver to the destination container and then by the transaction router to the request queue of a transaction executor. Once the sub-transaction is queued, the calling code gets a future back representing this computation. If the calling sub-transaction code does not synchronize

on the future, then once the caller completes, REACTDB enforces synchronization on the futures of all child sub-transactions. By the above synchronization policy, a root transaction can finish when all the sub-transactions created and invoked in its context finish, recursively. The transaction executor then invokes the transaction coordinator to initiate a commitment protocol across the containers that have been touched by the transaction, either directly or by any of its nested sub-transactions. The transaction coordinator in turn performs a 2PC protocol. The first phase of the protocol triggers validation of Silo’s OCC protocol on all the involved containers, during which locks are acquired on the write-set of the transaction. If any of the validations fail, the second phase of the protocol ensures that the transaction is aborted on all containers. Otherwise, the write phase of Silo’s OCC scheme is triggered. In either case, all locks are released appropriately [52].

3.2.3 Thread Management. To minimize the effect of stalls due to synchronization, each transaction executor maintains a thread pool to process (sub-)transactions. A configurable number of threads are allowed to become active and drain the transactor executor’s request queue, thus controlling the multi-programming level (MPL) per executor. In addition, the threads use cooperative multitasking to minimize context switching overheads. A thread blocks if it tries to access the result of a sub-transaction invoked on a different container and the result is not yet available. In such a situation, it notifies another thread to take over processing of the request queue and goes back to the thread pool when the (sub-)transaction being executed by it is completed.

3.3 Deployments

Configuration of transaction executors and containers allows infrastructure engineers to flexibly deploy REACTDB in a number of database architectures. In the remainder of the paper, we restrict ourselves to three main deployment strategies:

(S1) shared-everything-without-affinity: This strategy employs a single container in which each transaction executor can handle transactions on behalf of any reactor. REACTDB is configured with a round-robin router to load balance transactions among executors. All sub-transactions are executed within the same transaction executor to avoid any migration of control overhead. This strategy adheres to the architecture of most shared-everything databases [28]. **(S2) shared-everything-with-affinity:** This strategy is similar to shared-everything-without-affinity in that it employs a single container, but with the difference that an affinity-based router ensures that root transactions for a given reactor are processed by the same transaction executor. In sub-transaction calls, even if to different reactors, no migration of control happens, and the sub-transaction is executed by the same transaction executor of the root transaction. This deployment strategy closely adheres to the setup employed in the evaluation of Silo [52].

(S3) shared-nothing: This strategy employs as many containers as transaction executors, and a given reactor is mapped to exactly one transaction executor. While this strategy aims at maximizing program-to-data affinity, sub-transaction calls to different reactors may imply migration of control overheads to other transaction executors. We further decompose this configuration into shared-nothing-sync and shared-nothing-async, depending on how sub-transactions are invoked within application programs. In the former

option, sub-transactions are invoked synchronously by calling `get` on the sub-transaction’s future immediately after invocation. In the latter, the call to `get` is delayed as much as possible for maximal overlapping of application logic with sub-transaction calls. From an architecture perspective, both of these setups represent a shared-nothing deployment with differing application programs exercising different synchronization options. The shared-nothing-sync strategy models the setup of shared-nothing databases such as H-Store [50] and HyPer [30], albeit with a different concurrency control protocol due to potentially higher MPL per executor. The shared-nothing-async strategy allows REACTDB to further leverage intra-transaction parallelism as provided by the reactor programming model, exploiting cooperative multitasking.

Other flexible deployments, similar to [44], are possible as well. To change database architecture, only configuration files need to be edited and the system bootstrapped, without changes to application logic operating on reactors.

4 EVALUATION

In this section, we evaluate the effectiveness of REACTDB and the reactor programming model. The experiments broadly aim at validating the following hypotheses: **(H1)** The reactor programming model allows for reasoning about latency in alternative formulations of application programs (Section 4.2.1). **(H2)** The computational cost model of reactors can be efficiently realized by REACTDB, subject to its limitations (Section 4.2.2 and Appendix B). **(H3)** REACTDB allows for configuration of database architecture, without any changes to application code, so as to exploit asynchronicity in transactions depending on the level of load imposed on the database (Section 4.3 and Appendices C and D). We present additional evidence of REACTDB’s transactional scale-up capability in Appendix E. Furthermore, we evaluate procedure-level parallelism with reactors based on the scenario of Figure 1 in Appendix F.

4.1 Experimental Setup

Hardware. For our latency measurements in Section 4.2 and Appendix B, we employ a machine with one four-core, 3.6 GHz Intel Xeon E3-1276 processor with hyperthreading, leading to a total of eight hardware threads. Each physical core has a private 32 KB L1 cache and a private 256 KB L2 cache. All the cores share a last-level L3 cache of 8 MB. The machine has 32 GB of RAM and runs 64-bit Linux 4.1.2. A machine with high clock frequency and uniform memory access was chosen for these experiments to challenge our system’s ability to reflect low-level latency asymmetries in modern hardware as captured by our programming model.

For Section 4.3 and Appendices C–F, we use a machine with two sockets, each with 8-core 2.1 GHz AMD Opteron 6274 processors including two physical threads per core, leading to a total of 32 hardware threads. Each physical thread has a private 16 KB L1 data cache. Each physical core has a private 64 KB L1 instruction cache and a 2 MB L2 cache. Each of the two sockets has a 6 MB L3 cache. The machine has 125 GB of RAM in total, with half the memory attached to each of the two sockets, and runs 64-bit Linux 4.1.15. The higher number of hardware threads and accentuated cache coherence and cross-core synchronization effects allow us to demonstrate the effect of virtualization of database architecture in experiments varying transaction load and asynchrony.

Methodology. An epoch-based measurement approach similar to Oltpbench is used [21]. Average latency or throughput is calculated across 50 epochs and the standard deviation is plotted in error bars. All measurements include the time to generate transaction inputs.

Workloads and Deployments. For the experiments of Section 4.2, we implement an extended version of the Smallbank benchmark mix [3]. Smallbank simulates a banking application where customers access their savings and checking accounts. Oltpbench first extended this benchmark with a transfer transaction, which is implemented by a credit to a destination account and a debit from a source account [26]. We extend the benchmark further with a multi-transfer transaction. Multi-transfer simulates a group-based transfer, i.e., multiple transfers from the same source to multiple destinations. Thus, by varying the number of destination accounts for multi-transfer and controlling the deployment of REACTDB, we can vary both the amount of processing in the transaction as well as the amount of cross-reactor accesses that the transaction makes.

Each customer is modeled as a reactor. We configure REACTDB with 7 database containers, each hosting a single transaction executor for a total of 7 transaction executors mapped to 7 hardware threads. The deployment plan of REACTDB is configured so that each container holds a range of 1000 reactors. A single worker thread is employed to eliminate interference effects and allow us to measure latency overheads of single transactions. The worker thread generating transaction inputs and invocations is allocated in a separate worker container and pinned to the same physical core hosting the container responsible for the first range, but in a separate hardware thread. In order to keep our measurements comparable, the multi-transfer transaction input generator always chooses a source customer account from this first container.

The experiments of Section 4.3 use the classic TPC-C benchmark [51]. We closely follow the implementation of the benchmark from Oltpbench [26], which makes the usual simplifications, e.g., regarding think times. In our port of TPC-C, we model each warehouse as a reactor, and configure database containers differently according to the experiment. We vary the number of client worker threads generating transaction invocations, and group these workers into a worker container separate from the database containers that host transaction executors and carry out transaction logic. Each client worker thread generates load for only one warehouse (reactor), thus modeling client affinity to a warehouse. To showcase REACTDB’s ability to configure database architecture at deployment time, we experiment with the deployments described in Section 3.3.

Application Programs. We evaluate different application program formulations for the multi-transfer transaction added to Smallbank, exercising the asynchronous programming features of reactors. Similar to Figure 1(b), multi-transfer invokes multiple sub-transactions. In contrast to the figure, in some program variants, we force synchronous execution by immediately calling `get` on the future returned. The first formulation, *fully-sync*, invokes multiple transfer sub-transactions from the same source synchronously. Each transfer sub-transaction in turn invokes a synchronous credit sub-transaction on the destination account and a synchronous debit sub-transaction on the source account. The *partially-async* formulation behaves similarly; however, each transfer sub-transaction invokes an asynchronous credit on the destination account and a synchronous debit on the source account, overlapping half of the

writes in the processing logic while still executing communication proportional to the transaction size sequentially. The *fully-async* formulation does not invoke transfer sub-transactions, but rather explicitly invokes asynchronous credit sub-transactions on the destination accounts and multiple synchronous debit sub-transaction on the source account. Thus, not only are roughly half of the writes overlapped, but also a substantial part of the communication across reactors. The final formulation, *opt*, is similar to the *fully-async* transaction, but performs a single synchronous debit to the source account for the full amount instead of multiple debits. Consequently, processing depth is further reduced and should roughly equal two writes, while communication should be largely overlapped.

In addition, we implement all transactions of TPC-C in our programming model. Unless otherwise stated, we overlap calls between reactors as much as possible in transaction logic by invoking sub-transactions across reactors asynchronously.

4.2 Latency Control

4.2.1 Latency vs. Program Formulations. In this section, we show an experiment in which we vary the size of a multi-transfer transaction by increasing the number of destination accounts. Each destination is chosen on a different container out of the seven in our shared-nothing deployment. The latency for the different application program formulations is outlined in Figure 5. The observed curves match the trends predicted by the cost equation of Figure 3. First, as we increase transaction size, the processing and communication costs of a multi-transfer increase linearly across all formulations. Second, the highest latencies overall are for *fully-sync*, and latencies become lower as more asynchronicity is introduced in the formulations by overlapping sub-transaction execution. Third, there is a substantial gap between *partially-async* and *fully-async*, due to asymmetric costs between receiving procedure results and sending procedure invocations to other reactors. The latter manifests because of thread switching costs across cores in the receive code path, as opposed to atomic operations in the send code path. In *opt*, latency is further reduced when compared to *fully-async* by cutting the processing costs almost in half, which have a smaller impact than communication across cores. It is interesting to note that these optimizations can be done on the μsec scale. The programming model allows a developer to reduce the latency of a transaction from $86 \mu\text{sec}$ to $25 \mu\text{sec}$ by simple program reformulations without compromising consistency.

4.2.2 Cost Model Breakdown. In this section, we break down our measurements of transaction latencies by the cost components in Figure 3, and further validate that our cost model is realizable in REACTDB. We focus on the *fully-sync* and *opt* multi-transfer formulations described above, and vary the size of the multi-transfer transaction by changing the number of destination accounts similarly to Figure 5. For each variant, we profiled the execution time of the programs in REACTDB into the components of the cost model of Figure 3. In addition, we used the profiling information from *fully-sync* for a transaction size of one to calibrate the parameters of the cost model for prediction, including processing and communication costs. From the parameter values for the single-transfer *fully-sync* run, we employed the cost equation of Figure 3 to predict the execution costs for other transaction sizes and for both the

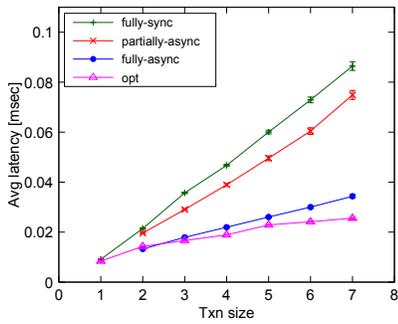


Figure 5: Latency vs. size and different user program formulations.

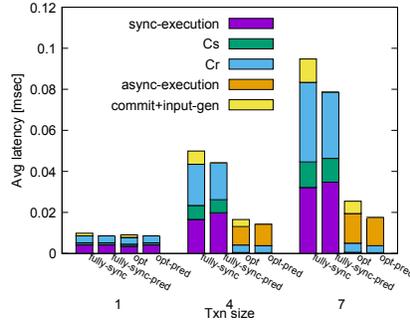


Figure 6: Latency breakdown into cost model components.

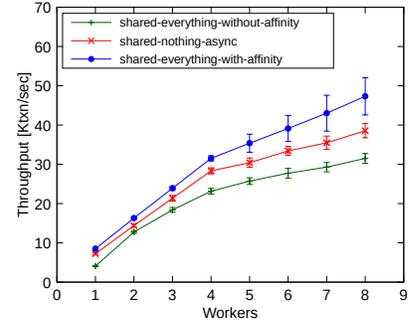


Figure 7: TPC-C throughput with varying load at scale factor 4.

fully-sync and opt program formulations. The predicted values are labeled fully-sync-pred and opt-pred.

We break down transaction latencies into the following components: (a) *sync-execution*: the cost of processing the logic in the transaction and in synchronous sub-transactions, corresponding to the first two components of the cost equation in Figure 3; (b) C_s and C_r : the forward and backward costs of communication between reactors in the third component of the cost equation; (c) *async-execution*: the cumulative execution cost of all asynchronous sub-transactions overlapped with any synchronous sub-transactions and processing logic, corresponding to the fourth component of the cost equation; (d) *commit + input-gen*: the cost of the commit protocol, including OCC and 2PC, along with the time to generate the inputs for the transaction. The latter cost component is not shown in Figure 3 since it only applies to root transactions and not to any sub-transaction in the reactor programming model. As such, we would expect the bulk of the difference between the predicted and observed performance to be explainable by this cost component.

Figure 6 shows that the predicted breakdown for the cost components closely matches the latencies profiled for actual executions in REACTDB, even at such a fine granularity. The slight difference in the overlap of different bars is within the variance of the observed vs. the calibration measurement, and expected especially since calibration measures parameters within the $5\mu\text{sec}$ range. For a transaction size of one, we can see that opt has the same performance behavior as fully-sync. This effect arises because the destination transaction executor for the credit in the transfer is the same as the source transaction executor for the debit, resulting in a synchronous execution of the credit and debit sub-transactions similar to fully-sync. As we increase the transaction size, the number of transaction executors spanned by the transaction increases, and the execution costs of fully-sync grow because of increasing costs in *sync-execution*, C_s and C_r . We again observe here cost asymmetry between C_s and C_r , arising for the same reasons remarked in Section 4.2.1. For opt, we do not observe any *sync-execution* costs, since all credit sub-transactions are overlapped with each other and with the single debit on the source reactor. The growth in the *async-execution* cost of opt with increasing transaction size is caused by the rising communication cost for the sequence of credit sub-transactions, i.e., the last asynchronous credit sub-transaction incurs a cumulative cost of communication of all asynchronous sub-transactions before it in the sequence.

In summary, we observe that the latencies of transactions can be reliably profiled in our system REACTDB into the components of the cost model of Figure 3 even in a challenging scenario where the cost of the processing logic in the benchmark is extremely small and comparable to the cost of communication across reactors. This evidence indicates that it is possible for developers to observe and explain the latency behavior of programs with reactors and to reformulate their programs for better performance.

4.3 Virtualization of Database Architecture

4.3.1 Effect of Load. We evaluate the behavior of the three database architecture deployments described in Section 3.3 with the standard mix of the TPC-C benchmark under increasing load from client workers. We fixed the database size to a scale factor of four, corresponding to four warehouse reactors. Correspondingly, we deployed four transaction executors in all configurations, but increased the number of workers invoking transactions on the database from 1 to 8. As such, we expect the database to get progressively more overloaded from 5 to 8 workers.

Figures 7 and 8 show the average throughput and latency of successful transactions in this scenario. We observe that shared-everything-with-affinity outperforms the other deployment configurations, since it exploits memory access affinities and minimizes cross-core communication overheads. Even though shared-nothing-async attempts to overlap cross-reactor sub-transactions as much as possible, the relative cost of dispatching sub-transactions with limited processing logic makes it more expensive than shared-everything-with-affinity, which employs direct memory accesses. We observe that the differences in throughput between the two methods are not large between 1 and 4 workers, since there is only a 1% chance of items for stock update in the new-order transaction and a 15% chance of customer lookups in the payment transaction being remote. Since workers employ client affinity to warehouses, note that for 5 workers, there are two worker threads generating workload for the first warehouse; for eight workers, there are two worker threads generating workload for every warehouse. So from 5 to 8 workers, the possibility of conflicts between concurrent transactions to the same warehouses arises, especially for the payment and new-order transactions.

As expected, therefore, from 1 to 4 workers, abort rates for all deployments are negligible and close to nil. Abort rates then go up to

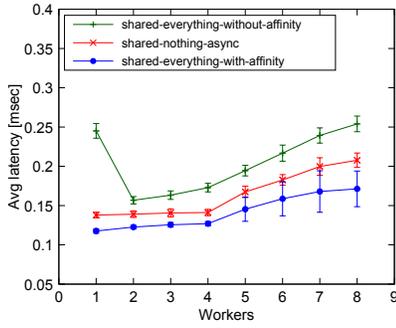


Figure 8: TPC-C latency with varying load at scale factor 4.

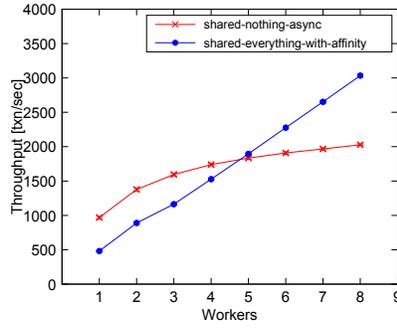


Figure 9: Throughput of new-order-delay transactions with varying load.

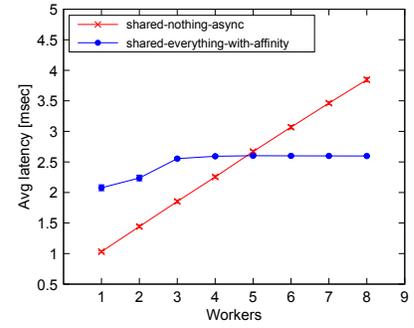


Figure 10: Latency of new-order-delay transactions with varying load.

4.1% for shared-everything-without-affinity and 5.72% for shared-nothing-async for eight workers. Interestingly, shared-everything-with-affinity is resilient to this effect, because each transaction executor runs a transaction to completion before picking up the next one from its queue. In other words, since shared-everything-with-affinity does not employ asynchronicity, transaction executor threads never block, and executors can operate with a multi-programming level of one. At the same time, affinity in routing ensures that each transaction executor services transactions for a different warehouse reactor, preserving identical execution behavior to that observed for workers 1 to 4 despite changes in the conflict patterns in workload generation. Even though shared-nothing-async also ensures affinity in transaction routing, it uses cooperative multi-threading and a higher multi-programming level for greater hardware utilization, and ends up executing multiple transactions concurrently from a given transaction executor’s queue.

While we increase the number of client workers, we must remember that the database system processing resources in the form of transaction executors remain fixed. So the throughput on average increases for all the deployments because the hardware utilization goes up. The hardware utilization on each of the transaction executor cores for shared-everything-with-affinity increases from 0 to 83% between 1 and 4 workers, and is pushed further to 99% as we get to eight workers. Given the use of asynchronicity in shared-nothing-async, this deployment uses all cores from the start, with a hardware utilization for one worker on the four transaction executor cores of 76%, 2.5%, 2.5% and 2.5%, respectively. The hardware utilization of transaction executor cores with this deployment is uniform at 79% with four workers, and keeps on rising up to 98% for eight workers. For shared-everything-without-affinity, the hardware utilization is uniform throughout, rising from 17% to 66% from 1 to 4 workers and reaching only 84% with eight workers.

We observe that shared-everything-without-affinity exhibits the worst performance. At one worker, every transaction invocation is routed by the round-robin load balancing router to a different transaction executor than the one that processed the last request, amplifying cross-core communication. As workers are added this effect diminishes; however, the lower hardware utilization and the eventual rise in abort rates limits the efficiency of this architecture.

In short, we remark that the capability to virtualize database architecture allows REACTDB to be configured to maximize hardware utilization and minimize conflicts in a standard OLTP benchmark.

Moreover, we observe that asynchronicity in transactions engenders a trade-off between communication overheads and processing costs. We validate this observation by fitting our cost model to the TPC-C new-order transaction in Appendix C, and explore asynchronicity trade-offs further in the next section.

4.3.2 Asynchronicity Tradeoffs. To additionally drill down on the potential benefits of asynchronicity under concurrency, we evaluate in this section the two database architectures shared-nothing-async and shared-everything-with-affinity under varying load. We control the amount of load by varying the number of workers from 1 to 8, while keeping the number of warehouses constant at a scale factor of eight. For clarity, we focus exclusively on new-order transactions. Each new-order consists of between 5-15 items and we force each of the items to be drawn from a remote warehouse with equal probability. Since the default new-order formulation has limited parallelism in the logic executed at remote warehouses, we augmented the logic for stock data update with an artificial delay between 300 and 400 μ sec by generating random numbers to model stock replenishment calculations. This increases the overall work in the transaction without increasing its data footprint and the contention on the database.

Figures 9 and 10 show the throughput and latency, respectively, of running 100% new-order-delay transactions under increasing load. With one worker, the throughput of shared-nothing-async is double that of shared-everything-with-affinity. The former executes all the stock updates across 5-6 remote warehouse asynchronously (average distinct remote warehouses chosen from 7 using a uniform distribution) fully utilizing the available hardware parallelism, while the latter executes the entire transaction logic sequentially. Although shared-nothing-async incurs higher communication cost in dispatching the stock updates to be performed by different warehouse reactors, the greater amount of work in each stock update makes it worthwhile in comparison to sequential shared memory accesses in shared-everything-with-affinity. Conversely, as we increase the number of workers and thus pressure on resources, the throughput of shared-nothing-async starts growing less than that of shared-everything-with-affinity. Note that the abort rate for the deployments was negligible (0.03-0.07%), highlighting the limited amount of contention on actual items.

In summary, these results suggest that the most effective database architecture may change depending on load conditions when asynchronicity can be exploited by transaction code. Under high load, shared-everything-with-affinity exhibits the best performance

among the architectures evaluated, since it reduces overhead at the expense of not utilizing at all intra-transaction parallelism. On the other hand, when load conditions are light to normal and when transaction logic comprises enough parallelism, shared-nothing-async can achieve substantially higher throughput and lower latency. To further validate these observations, we evaluate in Appendix D the effects of varying cross-reactor accesses in the TPC-C benchmark under conditions of high load.

5 RELATED WORK

In-memory OLTP Databases. H-Store [50] and HyPer [30] follow an extreme shared-nothing design by having single-threaded execution engines responsible for each data partition. As a result, single-partition transactions are extremely fast, but multi-partition transactions and skew greatly affect system throughput. LADS [56] improves upon this limitation by merging transaction logic and eliminating multi-partition synchronization through dynamic analysis of batches of specific transaction classes. In contrast to these shared-nothing engines, shared-everything lock-based OLTP systems specifically designed for multi-cores, such as DORA [42] and PLP [43], advocate partitioning of internal engine data structures for scalability. Orthrus [47] partitions only the lock manager and utilizes a message-passing design for lock acquisition to reduce lock contention across multiple cores for contended workloads.

In contrast to the baked-in architectural approach of earlier engines, REACTDB borrows the highly-scalable OCC implementation of Silo [52], building on top of it a virtualization layer that allows for flexible architectural deployments, e.g., as a classic shared-everything engine, a shared-nothing engine, or an affinity-based shared-everything engine. In addition, REACTDB is not restricted to specific transaction classes, supporting transactions with, e.g., user-defined aborts, conditionals, and range queries. Finally, REACTDB is the first engine realizing the programming model of reactors.

Transactional Partitioned Data Stores. A class of systems provides transactional support over key-value stores as long as keys are co-located in the same machine or key group [16, 17]. Warp [25], in contrast, provides full transaction support with nested transactions, but limits query capabilities, e.g., no predicate reads are provided nor relational query support. The limited transactional support and low-level storage-based programming model make it difficult to express applications as opposed to the reactor programming model, which provides serializable transactions with relational query capabilities. Recent work has also focused on enhancing concurrency through static analysis of transaction programs [38, 58]. The latter could be assimilated in the implementation of REACTDB’s concurrency control layers as future work.

Asynchronous Programming. As mentioned previously, reactors are a novel restructuring in the context of databases of the actor model [2]. In contrast to regular actors, reactors comprise an explicit memory model with transactions and relational querying, substantially simplifying program logic. These features make the reactor model differ significantly from the virtual actors of Orleans [7] and from other actor-based frameworks [4, 27]. Recent work in Orleans has focused on a vision of integrating traditional data-management functionality in a virtual actor runtime for the middle tier of a classic three-tier architecture [9, 23]. This approach is complementary to our work of integrating actor features in a

database system, i.e., enriching the data tier itself. Additionally, REACTDB comprises building a high-performance, scalable, multi-core OLTP system with an actor-oriented programming model and latency control, which is not the target design and feature set of the vision for Orleans [9].

As explained in Section 2.2, reactors are related to the early work on Argus [36] because of the asynchronous transactional programming model supporting nested function calls; however, the reactor programming model is substantially different from that of Argus. First, the use of a relational data and query model is a central idea of reactors, but not of Argus. Note that the latter is not a simple restriction of the former, because the programming issues handled by a relational abstraction, e.g., physical data independence, would need to be coded from scratch at a very low level in Argus. Second, user-defined logical actors are a central idea of reactors, but not of Argus. Guardians in Argus employ multiple concurrent processes explicitly, while reactors abstract away the mapping to threads in REACTDB. Third, reasoning about latency from the programming model is a central idea of reactors, but again not of Argus. Even though Argus has low-level asynchronous calls, it lacks an explicit cost model of synchronous and asynchronous communication. On the system implementation level, REACTDB is an OLTP database system designed for low-overhead virtualization of database architecture, which was never the focus of Argus. These differences to Argus also distinguish our work from a large class of object-oriented distributed computing and operating systems [11, 14, 18, 33, 41].

Database Virtualization. Virtualization of database engines for cloud computing has focused on particular target database architectures, e.g., shared-nothing databases with transactional support only within partitions [8] or distributed control architectures with weaker consistency guarantees [31]. By contrast, REACTDB offers infrastructure engineers the possibility to configure database architecture itself by containerization, while maintaining a high degree of transaction isolation. Our results support recent observations of low overhead of use of container mechanisms together with an in-memory database [39], while showing that even more flexibility in database architecture can be achieved at negligible cost.

The design of REACTDB is reminiscent of work on OLTP on hardware islands [44] and on the cost of synchronization primitives [19]. Complementary to recent work on compiler optimizations [46], reactors can be seen as a step towards improving programmability and support for stored procedures in database systems.

6 CONCLUSION

In this paper, we introduced reactors, a new relational abstraction for in-memory OLTP databases. Reactors comprise an asynchronous programming model that allows encoding of intra-transaction parallelism in database stored procedures while maintaining serializability. We presented the design of REACTDB, the first implementation of reactors. REACTDB enables flexible and controllable database architecture configuration at deployment time. Reactors open up a variety of directions for future work, ranging from reactor database modeling to efficient mapping of reactors to distributed hardware architectures.

Acknowledgments. We thank the anonymous reviewers for their insightful comments and Phil Bernstein for discussions and suggestions, which have helped us improve this paper.

REFERENCES

- [1] Alok Aggarwal and Jeffrey Scott Vitter. 1988. The Input/Output Complexity of Sorting and Related Problems. *Commun. ACM* 31, 9 (1988), 1116–1127.
- [2] Gul Agha. 1986. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, USA.
- [3] Mohammad Alomari, Michael J. Cahill, Alan Fekete, and Uwe Röhm. 2008. The Cost of Serializability on Platforms That Use Snapshot Isolation. In *Proc. ICDE*. 576–585.
- [4] Joe Armstrong. 2010. Erlang. *Commun. ACM* 53, 9 (2010), 68–75.
- [5] Subi Arumugam, Ravi Jampani, Luis Leopoldo Perez, Fei Xu, Christopher M. Jermaine, and Peter J. Haas. 2010. MCDB-R: Risk Analysis in the Database. *PVLDB* 3, 1 (2010), 782–793.
- [6] Catriel Beeri, Philip A. Bernstein, and Nathan Goodman. 1989. A Model for Concurrency in Nested Transactions Systems. *J. ACM* 36, 2 (1989), 230–269.
- [7] Philip Bernstein, Sergey Bykov, Alan Geller, Gabriel Kliot, and Jorgen Thelin. 2014. *Orleans: Distributed Virtual Actors for Programmability and Scalability*. Technical Report MSR-TR-2014-41. Microsoft Research.
- [8] Philip A. Bernstein, Istvan Cseri, Nishant Dani, Nigel Ellis, Ajay Kalhan, Gopal Kakivaya, David B. Lomet, Ramesh Manne, Lev Novik, and Tomas Talius. 2011. Adapting microsoft SQL server for cloud computing. In *Proc. ICDE*. 1255–1263.
- [9] Philip A. Bernstein, Mohammad Dashti, Tim Kiefer, and David Maier. 2017. Indexing in an Actor-Oriented Database. In *Proc. CIDR*.
- [10] Philip A. Bernstein, Vassos Hadzilacos, and Nathan Goodman. 1987. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley.
- [11] Kenneth P. Birman. 1985. Replication and Fault-Tolerance in the ISIS System. In *Proc. SOSP*. 79–86.
- [12] Andrew Brook. 2015. Low-latency distributed applications in finance. *Commun. ACM* 58, 7 (2015), 42–50.
- [13] Michael J. Carey, David J. DeWitt, and Scott L. Vandenberg. 1988. A Data Model and Query Language for EXODUS. In *Proc. ACM SIGMOD*. 413–423.
- [14] Panos K. Chrysanthis, Krithi Ramamritham, David W. Stemple, and Stephen Vinter. 1986. The Gutenberg Operating System Kernel. In *Proc. FJCC*. 1159–1167.
- [15] Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghuram Ramakrishnan, and Russell Sears. 2010. Benchmarking cloud serving systems with YCSB. In *Proc. ACM SoCC*. 143–154.
- [16] Sudipto Das, Divyakant Agrawal, and Amr El Abbadi. 2010. G-Store: a scalable data store for transactional multi key access in the cloud. In *Proc. ACM SoCC*. 163–174.
- [17] Sudipto Das, Amr El Abbadi, and Divyakant Agrawal. 2009. ElasTraS: An Elastic Transactional Data Store in the Cloud. In *Workshop on Hot Topics in Cloud Computing, HotCloud*.
- [18] Partha Dasgupta, Richard J. LeBlanc, and William F. Appelbe. 1988. The Clouds Distributed Operating System. In *Proc. ICDCS*. 2–9.
- [19] Tudor David, Rachid Guerraoui, and Vasileios Trigonakis. 2013. Everything You Always Wanted to Know About Synchronization but Were Afraid to Ask. In *Proc. SOSP*. 33–48.
- [20] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Ake Larson, Pravin Mittal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. 2013. Hekaton: SQL Server’s Memory-optimized OLTP Engine. In *Proc. ACM SIGMOD*. 1243–1254.
- [21] Djellel Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudré-Mauroux. 2013. OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases. *PVLDB* 7, 4 (2013), 277–288.
- [22] Aleksandar Dragojević, Dushyanth Narayanan, Edmund B. Nightingale, Matthew Renzelmann, Alex Shamis, Anirudh Badam, and Miguel Castro. 2015. No Compromises: Distributed Transactions with Consistency, Availability, and Performance. In *Proc. SOSP*. 54–70.
- [23] Tamer Eldeeb and Phil Bernstein. 2016. *Transactions for Distributed Actors in the Cloud*. Technical Report MSR-TR-2016-1001. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/transactions-distributed-actors-cloud-2/>
- [24] E. N. Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. 2002. A survey of rollback-recovery protocols in message-passing systems. *ACM Comput. Surv.* 34, 3 (2002), 375–408.
- [25] Robert Escriva, Bernard Wong, and Emin Gün Sirer. 2015. Warp: Lightweight Multi-Key Transactions for Key-Value Stores. *CoRR* abs/1509.07815 (2015).
- [26] H-Store 2017. H-Store Supported Benchmarks (Smallbank). (May 2017). <http://hstore.cs.brown.edu/documentation/deployment/benchmarks/>
- [27] Yaroslav Hayduk, Anita Sobe, and Pascal Felber. 2015. Dynamic Message Processing and Transactional Memory in the Actor Model. In *Proc. DAIS*. 94–107.
- [28] Joseph Hellerstein, Michael Stonebraker, and James Hamilton. 2007. Architecture of a Database System. *Foundations and Trends in Databases* 1, 2 (2007), 141–259.
- [29] Jiamin Huang, Barzan Mozafari, Grant Schoenebeck, and Thomas F. Wenisch. 2017. A Top-Down Approach to Achieving Performance Predictability in Database Systems. In *Proc. ACM SIGMOD*. 745–758.
- [30] Alfons Kemper and Thomas Neumann. 2011. HyPer: A Hybrid OLTP&OLAP Main Memory Database System Based on Virtual Memory Snapshots. In *Proc. ICDE*. 195–206.
- [31] Donald Kossmann, Tim Kraska, and Simon Loesing. 2010. An evaluation of alternative architectures for transaction processing in the cloud. In *Proc. ACM SIGMOD*. 579–590.
- [32] Per-Ake Larson (Ed.). 2013. *IEEE Data Engineering Bulletin: Special Issue on Main-Memory Database Systems*. Vol. 36, number 2.
- [33] Edward D. Lazowska, Henry M. Levy, Guy T. Almes, Michael J. Fischer, Robert J. Fowler, and Stephen C. Vestal. 1981. The Architecture of the Eden System. In *Proc. SOSP*. 148–159.
- [34] Christophe Lécluse and Philippe Richard. 1989. The O2 Database Programming Language. In *Proc. VLDB*. 411–422.
- [35] Daan Leijen, Wolfram Schulte, and Sebastian Burckhardt. 2009. The design of a task parallel library. In *Proc. OOPSLA*. 227–242.
- [36] Barbara Liskov, Dorothy Curtis, Paul Johnson, and Robert Scheifler. 1987. Implementation of Argus. In *Proc. SOSP*. 111–122.
- [37] B. Liskov and L. Shrira. 1988. Promises: Linguistic Support for Efficient Asynchronous Procedure Calls in Distributed Systems. In *Proc. PLDI*. 260–267.
- [38] Shuai Mu, Yang Cui, Yang Zhang, Wyatt Lloyd, and Jinyang Li. 2014. Extracting More Concurrency from Distributed Transactions. In *Proc. OSDI*. 479–494.
- [39] Tobias Mühlbauer, Wolf Rödiger, Andreas Kipf, Alfons Kemper, and Thomas Neumann. 2015. High-Performance Main-Memory Database Systems and Modern Virtualization: Friends or Foes?. In *Proc. DanC*. 4:1–4:4.
- [40] Vivek R. Narasayya, Sudipto Das, Manoj Syamala, Badrish Chandramouli, and Surajit Chaudhuri. 2013. SQLVM: Performance Isolation in Multi-Tenant Relational Database-as-a-Service. In *Proc. CIDR*.
- [41] Howard Oakley. 1989. Mercury: an operating system for medium-grained parallelism. *Microprocessors and Microsystems - Embedded Hardware Design* 13, 2 (1989), 97–102.
- [42] Ippokratis Pandis, Ryan Johnson, Nikos Hardavellas, and Anastasia Ailamaki. 2010. Data-oriented Transaction Execution. *PVLDB* 3, 1-2 (2010), 928–939.
- [43] Ippokratis Pandis, Pinar Tözün, Ryan Johnson, and Anastasia Ailamaki. 2011. PLP: Page Latch-free Shared-everything OLTP. *PVLDB* 4, 10 (2011), 610–621.
- [44] Danica Porobic, Ippokratis Pandis, Miguel Branco, Pinar Tözün, and Anastasia Ailamaki. 2012. OLTP on Hardware Islands. *PVLDB* 5, 11 (2012), 1447–1458.
- [45] PostgreSQL Documentation 2017. Parallel Safety in PostgreSQL 10 (Section 15.4). (October 2017). <https://www.postgresql.org/docs/10/static/parallel-safety.html>
- [46] Karthik Ramachandra, Kwanghyun Park, K. Venkatesh Emani, Alan Halverson, César A. Galindo-Legaria, and Conor Cunningham. 2017. Optimization of Imperative Programs in a Relational Database. *PVLDB* 11, 4 (2017), 432–444.
- [47] Kun Ren, Jose M. Faleiro, and Daniel J. Abadi. 2016. Design Principles for Scaling Multi-core OLTP Under High Contention. In *Proc. ACM SIGMOD*. 1583–1598.
- [48] R. Shoup and D. Pritchett. 2006. The eBay architecture. In *SD Forum*.
- [49] Michael Stonebraker. 2012. New opportunities for New SQL. *Commun. ACM* 55, 11 (2012), 10–11.
- [50] Michael Stonebraker, Samuel Madden, Daniel J. Abadi, Stavros Harizopoulos, Nabil Hachem, and Pat Helland. 2007. The End of an Architectural Era: (It’s Time for a Complete Rewrite). In *Proc. VLDB*. 1150–1160.
- [51] TPCC 2017. The TPC-C Benchmark. (May 2017). <http://www.tpc.org/tpcc/>
- [52] Stephen Tu, Wenting Zheng, Eddie Kohler, Barbara Liskov, and Samuel Madden. 2013. Speedy Transactions in Multicore In-memory Databases. In *Proc. SOSP*. 18–32.
- [53] Pattishall D. V. 2007. Federation at Flickr (doing billions of queries per day). In *MySQL Conference*.
- [54] Gerhard Weikum and Gottfried Vossen. 2002. *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery*. Morgan Kaufmann.
- [55] Walker M. White, Christoph Koch, Nitin Gupta, Johannes Gehrke, and Alan J. Demers. 2007. Database research opportunities in computer games. *SIGMOD Record* 36, 3 (2007), 7–13.
- [56] Chang Yao, Divyakant Agrawal, Gang Chen, Qian Lin, Beng Chin Ooi, Weng-Fai Wong, and Meihui Zhang. 2016. Exploiting Single-Threaded Model in Multi-Core In-Memory Systems. *IEEE Trans. Knowl. Data Eng.* 28, 10 (2016), 2635–2650.
- [57] Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, and Michael Stonebraker. 2014. Staring into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores. *PVLDB* 8, 3 (2014), 209–220.
- [58] Yang Zhang, Russell Power, Siyuan Zhou, Yair Sovran, Marcos K. Aguilera, and Jinyang Li. 2013. Transaction chains: achieving serializability with low latency in geo-distributed storage systems. In *Proc. SOSP*. 276–291.
- [59] Wenting Zheng, Stephen Tu, Eddie Kohler, and Barbara Liskov. 2014. Fast Databases with Fast Durability and Recovery Through Multicore Parallelism. In *Proc. OSDI*. 465–477.

A PROOF OF THEOREM 2.7

Let us assume H is serializable and H' is not serializable. From the serializability theorem, since H is serializable, the serializability graph of H ($SG(H)$) is acyclic; since the projected history H' is not serializable, the serializability graph $SG(H')$ must be cyclic.

Therefore, there must exist a cycle $T'_i \rightarrow \dots \rightarrow T'_j \rightarrow \dots \rightarrow T'_i$. Since the graph is built on operations of the classic transactional model, then there must be conflicting operations $o'_i <_{P_H} \dots <_{P_H} o'_j <_{P_H} \dots <_{P_H} o'_i$. By condition (3) of Definition 2.6, there must exist sub-transactions $ST_{i,l}^k \in T_i$ and $ST_{j,l'}^{k'} \in T_j$ such that $ST_{i,l}^k <_H \dots <_H ST_{j,l'}^{k'} <_H \dots <_H ST_{i,l}^k$. As a result, $SG(H)$ must be cyclic, and we arrive at a contradiction. To show the reverse direction, it is simple to follow a similar argument, but starting with a cyclic graph in $SG(H)$ and showing that $SG(H')$ must be cyclic as well in contradiction.

B EFFECT OF SKEW AND QUEUING ON LATENCY

In this appendix, we drill down experimentally into the limitations discussed for our cost model in Section 2.4, evaluating a scenario with queuing delays and skew.

Setup. We created a separate setup to evaluate limitations in explanatory ability of our cost model by augmenting the YCSB benchmark [15] with a `multi_update` transaction that updates 10 keys. We model each key as a reactor, and the `multi_update` transaction invokes a read-modify-write update sub-transaction asynchronously for each key (record size of 100 bytes). Unlike in the setup with Smallbank, we configure REACTDB such that transactions may not realize full parallelism at the system level. In particular, we chose a scale factor of 4 where each scale factor corresponds to 10,000 keys. Four containers are deployed, each with one transaction executor, and assigned 10,000 contiguous reactors corresponding to the scale factor. Furthermore, we select the reactor where the `multi_update` transaction is invoked randomly from the set of 10 keys in the update, and choose the keys for `multi_update` from a zipfian distribution. As such, we model program logic where it is not possible to determine statically how many sub-transactions will be realized by the system synchronously or asynchronously. However, to ensure that transactions remain fork-join, we sorted the keys to be updated in a `multi_update` transaction such that keys deployed in remote transaction executors precede the ones in the local transaction executor where the transaction is initiated. Finally, the experiment is performed both with one and four workers to simulate absence and presence of queuing delays, respectively.

Results. We employ the `multi_update` transaction under a 100% mix. Recall that our setup explained above explicitly restricts the parallelism that can be realized by the system. So to be able to fit our cost model, we recorded the average sizes of the realized $sync_{ovp}(ST_{i,j}^k)$ and $async(ST_{i,j}^k)$ sequences with the zipfian distribution employed. Similar to the experiment in Section 4.2.2, we calibrated the average communication and processing cost parameters by profiling the `multi_update` transaction with updates to a single key chosen from a uniform distribution. We emphasize that the cost model abstracts intra-transaction parallelism to aid developers in contrasting transaction program formulations, and thus does not capture interference among concurrent workers or queuing effects. As such, we expect cost model predictions to be most accurate in configurations where a single worker is deployed, as in Section 4.2.

Figure 11 shows the observed behavior of average latency with varying skew, as captured by the zipfian constant used to select

Cross Reactor Access %	1 worker				4 workers	
	TPS	Latency msec			TPS	Latency msec
	Obs	Pred	Pred+C+I	Obs	Obs	Obs
1	6921	0.131	0.148	0.144	27091	0.148
100	5246	0.159	0.189	0.191	14485	0.277

Table 1: TPC-C new-order performance at scale factor 4.

keys for `multi_update`, for different number of workers along with the cost model predictions. For one worker, the observed average latency decreases as skew increases, given that more of the sub-transactions invoked in `multi_update` become synchronous. This effect arises since the communication overhead to dispatch a sub-transaction to a remote reactor is greater than the time to process a single update. This trend is also captured in our cost model prediction, 1 worker pred, which shows decreasing average latency as the zipfian constant increases to 0.99. Interestingly, the curve predicts that the average latency should increase after 0.99, which conflicts with our observation. This is because the cost of processing actually decreases when all updates are to the same key in the transaction. In this case, the read in a read-modify-write can just return the previous updated value from the transaction’s write-set instead of accessing the index structure, which was not accounted for while calibrating processing cost. In addition to 1 worker pred, we also show a curve adding to the prediction the measured costs of input generation and commitment. We remind the reader that the latter two costs are not part of the equation in Figure 3, and as in Section 4.2.2, the difference between predicted and observed latencies for a single worker are largely accounted by these two factors.

With four workers, queuing and skew together lead to both higher and more variable latencies overall, which are as expected not captured by the cost model. Abort rates also increase, going from 0.26% to 3.24% as the zipfian constant increases from 0.01 to 0.99. Hardware utilization on the transaction executor core hosting the most accessed reactors reflects the trend in the figure, corresponding to 63%, 92% and 100% at zipfian constants 0.01, 0.5 and 0.99. At 5.0 skew, a single reactor is accessed, eliminating variability but not queuing.

C COST MODEL VALIDATION WITH TPC-C NEW-ORDER

We observed in Section 4.3.1 that shared-nothing-async underperforms compared to shared-everything-with-affinity even with a single client worker, despite having more transaction executors available for transaction processing. To better understand this effect, we turned to our cost model for validation of a mix consisting of 100% new-order transactions with four warehouses and four transaction executors. Similarly to what was done in Section 4.2.2 and Appendix B, we calibrated cost model parameters with a new-order transaction requesting only one item from a local and one from a remote warehouse. Moreover, we recorded the average numbers of synchronous and asynchronous stock-update requests realized with a single worker, as well as the average numbers of items involved. To explore the impact of communication overheads, we evaluated two extremes with respectively 1% and 100% probability of cross-reactor accesses in stock updates (see also Appendix D).

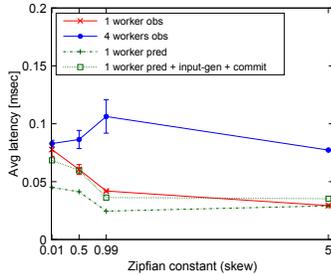


Figure 11: Effect of skew and queuing on latency.

Table 1 summarizes the results obtained with one and four workers. We obtain excellent fit between the cost model prediction after including commit and input generation cost (Pred+C+I) and the observed latency for one worker under both 1% and 100% cross-reactor accesses. We observed that the cost of communication between reactors, especially C_r as remarked in Section 4.2.1, is high compared with the processing cost of stock updates. However, the relatively small growth in the latency for 100% cross-reactor access with one worker compared with 1% cross-reactor accesses shows the benefits of overlapping multiple asynchronous sub-transactions across reactors. With four workers, queueing effects manifest with 100% cross-reactor accesses, which are as discussed in Appendix B not part of the cost model, though prediction is still accurate for 1% cross-reactor accesses.

The hardware utilization with one worker on the four transaction executor cores under 1% cross-reactor accesses is 78%, 2.3%, 2.3%, and 2.3%, respectively, while the values under 100% cross-reactor accesses are 65%, 24%, 24%, and 24%. The use of four workers increases utilization, which becomes uniform at 82% and 86% under 1% and 100% cross-reactor accesses, respectively.

D EFFECT OF CROSS-REACTOR TRANSACTIONS

In this appendix, we evaluate the impact of cross-reactor transactions on the performance of the different database architectures, complementing the results of Section 4.3. For clarity, we focus exclusively on new-order transactions without any additional computations. To vary the percentage of cross-reactor new-order transactions, we vary the probability that a single item in the transaction is drawn from a remote warehouse (the remote warehouses are again chosen with an equal probability). Each new-order consists of between 5-15 items.

Figure 12 shows the average latencies of running the TPC-C benchmark with 100% new-order transactions at a scale factor of eight, i.e., eight warehouses receive transactions from eight workers at peak load. Due to space constraints, we omit measured throughputs. Since REACTDB uses Silo’s OCC protocol and owing to the low contention in TPC-C even upon increasing the number of remote items, we would expect the latency of shared-everything-with-affinity to be agnostic to changes in the proportion of cross-reactor transaction as per the results in [52]. However, we see a small gradual increase in the latency for all the deployments. We believe these effects are a consequence of cache coherence and cross-core communication overheads. We observe further that both shared-nothing-sync and shared-nothing-async configurations exhibit the

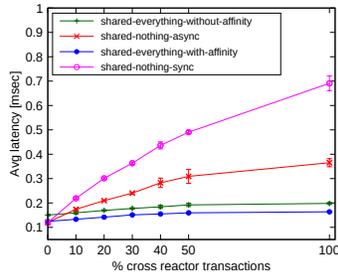


Figure 12: Latency of cross-reactor TPC-C new-order (scale factor 8).

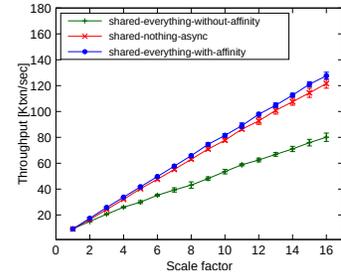


Figure 13: TPC-C throughput with varying deployments.

same latency at 0% cross-reactor transactions as shared-everything-with-affinity. However, there is a sharp drop in the performance of shared-nothing deployments from 0% to 10% cross-reactor transactions. This effect is in line with our previous observation that sub-transaction invocations require expensive migration of control in contrast to both shared-everything-without-affinity and shared-everything-with-affinity. Note that the abort rate for all deployments remained negligible (0.02%-0.04%), highlighting the limited amount of contention on actual items.

We observe that shared-nothing-async exhibits higher resilience to increase in cross-reactor transactions when compared with shared-nothing-sync. In particular, latency of shared-nothing-async is better by roughly a factor of two at 100% cross-reactor transactions. This is because shared-nothing-async employs new-order transactions with asynchronous sub-transaction invocations on remote warehouse reactors, and tries to overlap remote sub-transaction invocation with execution of logic locally on a warehouse reactor. This demonstrates how application programs can leverage the programming model to engineer application code using reactors with different performance characteristics. At the same time, infrastructure engineers can select the database architecture that best fits the execution conditions for the workload without changes to application code. In the case of peak load and limited intra-transaction parallelism, shared-everything-with-affinity turned out to be the best architecture among the ones considered for this scenario, in line with the results of [52].

E SCALE-UP AND OVERHEAD IN REACTDB

This appendix further complements Section 4.3 by presenting transactional scale-up in REACTDB to the degree allowed by our hardware setup and discussing containerization overhead.

Transactional Scale-Up. In this section, we evaluate the scalability of REACTDB across multiple cores for the three database architecture deployments described in Section 3.3. Figure 13 shows the average transaction throughput of running the TPC-C transaction mix as we increase the number of warehouses (reactors). While we have also measured latencies, we omit the results due to space constraints. Note that we configure the number of transaction executors to be equal to the scale factor for the experiment.

We observe that the shared-everything-without-affinity deployment exhibits the worst throughput scalability among the deployments selected. This effect is a consequence of shared-everything-without-affinity’s poor ability to exploit memory access affinities within each transaction executor, given round-robin routing of transactions. On the other hand, shared-everything-with-affinity

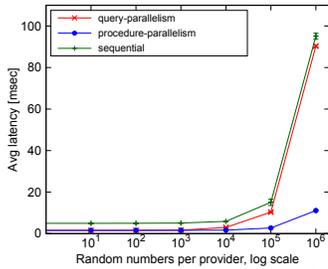


Figure 14: Latency of query- vs. procedure-level parallelism.

and shared-nothing-async both take advantage of access affinities and behave similarly. We see that shared-everything-with-affinity is slightly superior to shared-nothing-async. The difference lies in the relative costs in these deployments of sub-transaction invocations vs. direct memory accesses of data for remote warehouses. For a scale factor of one, there are no cross-reactor transactions, and the performance of the two deployments is identical. From a scale factor of two onwards, the probabilities of cross-reactor transactions range between 0% to 10%, as discussed in Section 4.3.1. In shared-nothing-async, a sub-transaction call is routed to its corresponding transaction executor, incurring context switching and communication overheads. By contrast, since shared-everything-with-affinity executes the sub-transaction in the same transaction executor, the remote call costs are traded off for the relatively smaller costs of cache pressure. We also ran the experiment with all the transaction classes in the TPC-C mix invoking sub-transactions synchronously in the shared-nothing deployment (shared-nothing-sync configuration described in Section 3.3). However, the throughput and latency of this configuration was close (within the variance bars) to the shared-nothing-async configuration because of the low percentage of remote warehouse calls in the default TPC-C mix. We hence omit the curve from Figure 13 for brevity.

In short, the results indicate that REACTDB can be flexibly configured with different database architectures to achieve adequate transactional scalability in our available hardware setup for the standard TPC-C workload. Further, high affinity of data accesses to physical processing elements (cores) is crucial to performance.

Containerization Overheads. To account for the overhead of containerization, we also ran REACTDB while submitting empty transactions with concurrency control disabled. We observe roughly constant overhead per transaction invocation across scale factors of around 22 μ sec. We measured that thread switching overhead between the worker and transaction executor across different cores is a largely dominant factor and is dependent on the machine used. The overhead accounted for 18% of average TPC-C transaction latency. When compared with executing the TPC-C application code directly within the database kernel without worker threads generating transaction invocations that are separate from database processing threads, as in Silo, the overhead is significant, but if a database engine with kernel thread separation is assumed, as is the normal case, the overhead is negligible.

F PROCEDURE-LEVEL PARALLELISM WITH REACTORS

In this section, we revisit the example introduced in Figure 1 and evaluate the potential performance gains that can be obtained by

leveraging procedure-level parallelism in the reactor model. In Figure 1(a), we see the formulation of `auth_pay` in the classic transactional model. If the `orders` relation were partitioned, the join query between `provider` and `orders` would be eligible for parallelization by a query optimizer in a traditional relational database. Note that the optimizer would consider query-level parallelization, and not a holistic procedure-level parallelization achievable by manual decomposition of the join in the reactor programming model as shown in Figure 1(b). In the latter, the potentially expensive computation of `sim_risk` would also be parallelized across reactors.

For this experiment, we configured REACTDB with up to 16 transaction executors. To focus on intra-transaction parallelism, a single worker generates `auth_pay` transaction invocations targeted at one Exchange and 15 Provider reactors, used to express three program execution strategies in REACTDB. Both the sequential and query-parallelism strategies encode the program of Figure 1(a) for the classic transactional model. In sequential, we deploy a single container and transaction executor for all reactors in REACTDB, and thus the whole procedure body executes in a single hardware thread. In query-parallelism, by contrast, we horizontally decompose `orders` across the Provider reactors by deploying one container and transaction executor for each reactor. Similar to a partitioned database, this results in a parallel foreign key join between `providers` and `orders` fragments, but sequential execution of the remainder of the procedure. Finally, the procedure-parallelism strategy encodes the program of Figure 1(b). Using the same deployment as for query-parallelism, this achieves holistic procedure parallelization in the reactor model. The inputs for `auth_pay` were chosen using uniform distributions, and the `orders` relation was loaded with 30,000 records per provider, mirroring the cardinality of the `orders` relation in the TPC-C benchmark. To simulate another transaction settling orders and keep the number of records scanned in the `orders` relation fixed, a pre-configured window of records is reverse range scanned ordered by time per provider. This window parameter has a direct effect on the benefit of query-parallelism compared with sequential. For clarity, we tuned this value to 800 records per provider to ensure that query-parallelism would outperform sequential by 4x when `sim_risk` is not invoked. We simulated the computational load of `sim_risk` by random number generation similar to Section 4.3.2. We also loaded our data values such that `sim_risk` is always invoked and transactions are not aborted due to application logic failures.

Figure 14 shows average latencies as we vary the computational load in `sim_risk` by generating progressively more random numbers. We observe that procedure-parallelism is more resilient to changes in computational load and better utilizes the available parallel hardware in this benchmark. At the extreme of 10^6 random numbers generated per provider, the latency of procedure-parallelism is a factor of 8.14x and 8.57x lower than that of query-parallelism and sequential, respectively. The transaction executor core hosting the Exchange reactor becomes overloaded at 100% utilization under query-parallelism, while with procedure-parallelism utilization remains uniform at 84% across executor cores hosting Provider reactors and 14% in the Exchange executor core. In sequential, the transaction executor mapped to all reactors becomes overloaded at 100% even when the computational load is zero.