# P4DB - The Case for In-Network OLTP

Matthias Jasny Technical University of Darmstadt Lasse Thostrup Technical University of Darmstadt

# ABSTRACT

In this paper we present a new approach for distributed DBMSs called P4DB, that uses a programmable switch to accelerate OLTP workloads. The main idea of P4DB is that it implements a transaction processing engine on top of a P4-programmable switch. The switch can thus act as an accelerator in the network, especially when it is used to store and process hot (contended) tuples on the switch. In our experiments, we show that P4DB hence provides significant benefits compared to traditional DBMS architectures and can achieve a speedup of up to 8×.

#### **CCS CONCEPTS**

Information systems → Distributed database transactions;
 Networks → In-network processing.

# **KEYWORDS**

Distributed Databases, Transactions, Programmable Switches, P4

#### **ACM Reference Format:**

Matthias Jasny, Lasse Thostrup, Tobias Ziegler, and Carsten Binnig. 2022. P4DB - The Case for In-Network OLTP. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22), June 12–17,* 2022, Philadelphia, PA, USA. ACM, New York, NY, USA, 15 pages. https: //doi.org/10.1145/3514221.3517825

#### **1 INTRODUCTION**

*Motivation.* The efficient use of data center networks plays a significant role on the performance of distributed DBMSs. Traditionally, distributed DBMSs were built on the assumption that the network is a major bottleneck. As such, classical distributed DBMSs were designed to mitigate the effects of the high network cost using sophisticated techniques such as complicated partitioning schemes [15, 48, 52, 70], semi-join transformations [47, 50, 53], speculative execution [49], new consistency levels [35], or even the relaxation of atomicity guarantees [37, 40]. However, data center networks have been evolving significantly in recent years.

A first major trend that we have seen in the last years is that data center networks have evolved *from being slow to being fast*. For example, when looking at the network speed provided by cloud vendors such as Amazon, Microsoft or Google for their hosted cloud instances, we see that even for the smaller (i.e., cost-efficient) instances, the network link has a speed of at least 10Gbps and can

SIGMOD '22, June 12-17, 2022, Philadelphia, PA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9249-5/22/06...\$15.00 https://doi.org/10.1145/3514221.3517825



Figure 1: OLTP processing in P4DB using a programmable switch. Compared to a traditional DBMS without using the switch, P4DB provides significant speedups. Details about the setup are explained in our evaluation.

reach up to 100Gbps for larger instance types. Consequently, there are significant efforts in database community into investigating how these high-speed networks can be used efficiently by redesigning distributed DBMSs for zero-copy protocols such as RDMA or DPDK to avoid the overhead of classical network stacks such as TCP/IP [7, 30, 64, 69, 72, 74].

Another major trend that we have seen in the recent years is that with the rise of software-defined networks [36], switches and network cards have become programmable [10, 51] and thus turned networks *from being passive to being active*. This programmability of the network opens up many additional opportunities to tailor the network to the applications on top. In particular, the programmability of the so-called data plane allows applications to offload computation to the network devices (aka in-network-processing or INP for short). INP has shown to provide significant benefits for distributed data processing in general, including key-value stores [28, 41, 68] as well as distributed OLAP and ML [8, 39, 42, 57, 61]. However, as far as OLTP is concerned, there has been only limited work so far which offloads only certain sub-components of OLTP, such as lock management [68] or replication protocols [28, 73].

*Contributions.* As a main contribution, in this paper we present a new, more aggressive approach for distributed DBMSs called P4DB that involves a programmable switch more actively to accelerate OLTP workloads. The main idea of P4DB is that it implements a *full transaction processing engine on top of a programmable switch*, using P4 as the de facto data plane programming language. That way, P4DB exposes the switch as an "additional" database node that, however, comes with very different properties compared to a normal database node (as we discuss below). As a result, the switch is an ideal "place" to store and process hot (contended) tuples that typically lead to a significant performance degradation in traditional DBMSs. As we show in Figure 1, P4DB can thus significantly speed up OLTP processing in distributed DBMSs for various workloads in case hot tuples are offloaded to the switch.

To better understand where the benefits of using a switch originate from, let us first look at the hardware characteristics of (programmable) switches that make them an interesting candidate for hot tuples. As a main difference to a normal database node of the host DBMS, a programmable switch comes with two interesting characteristics: (1) Programmable switches are designed to process

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 the author(s) 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.

the aggregated load of network traffic from all connected servers at line-rate. (2) The switch can be reached from database nodes that are directly connected to the switch with only half of the network latency compared to the latency required to reach any other (remote) database node that is connected to the same switch. Based on these characteristics, we next argue that a programmable switch is an ideal place for storing and processing hot items of an OLTP workload that are most frequently accessed by transactions.

In a traditional distributed DBMS architecture, hot items that are accessed frequently by transactions typically lead to a severely degraded performance, as mentioned before. This is especially true, when hot items are involved in distributed transactions, as hot items typically suffer from increased remote access latencies. The latency increase, in turn, leads to an increased likelihood of contention on those items and thus higher abort rates [71]. In contrast, P4DB can mitigate these effects if hot items are stored and processed on a programmable switch: This is because when hot tuples are involved in distributed transactions, storing them on a switch reduces access latencies, which reduces the overall contention, directly leading to performance benefits.

Surprisingly, however, as we show in our evaluation, even if an OLTP workload is perfectly partitionable and no distributed transactions are involved in a workload, or even if there is only limited skew in a workload, P4DB can still speed up processing compared to a traditional DBMS. The reason is that programmable switches provide an execution model that enables a new lock-free execution scheme for transactions, which allows the switch in P4DB to process transactions on hot tuples at high speeds since it avoids any contention. However, implementing a transaction engine on top of a programmable switch does not come "for free" and many challenges need to be addressed. For example, typical programmable switches of today, such as recent Tofino-based switches [5], have restrictions not only on the memory model (i.e., how data in the switch memory can be accessed by in-switch-programs), but also what kinds of in-switch-programs are supported by their execution models. By making clever use of the switch memory and the execution models of programmable switches though, we can efficiently support the concurrent execution of transactions in an abort-free manner on the switch, as mentioned before.

Finally, we think P4DB is deployable and compatible with existing datacenter networks. As a main deployment model, P4DB targets cloud providers that have dedicated racks for database services, which is a common scheme to provide high performance for distributed DBMSs [68]. For such a deployment, P4DB only needs to augment the Top-of-Rack (ToR) switches with a custom data plane module for processing transactions on hot tuples. But other deployments, e.g. hierarchies of switches or multi-tenant deployments, are clearly also interesting. While such deployments are out of the scope of this paper and are important avenues of future work, at the end of the paper we provide a short discussion of how we think P4DB can generalize to such deployments.

*Outline.* The remainder of this paper is organized as follows: First, in Section 2 we discuss the relevant background on programmable switches and their memory and execution models. In Section 3 we provide an overview of P4DB before we then explain in Section 4 and Section 5 the details of the storage and execution model we implemented to execute database transactions on



Figure 2: The Protocol Independent Switch Architecture (PISA): All stages (Parser, MAU Stages and Deparser) are programmable and allow flexible packet processing in the data plane based on packet headers and metadata.

a programmable switch. Afterwards, in Section 6 we then discuss the integration of the in-switch transaction processing into a host DBMS. Finally, we conclude with an extensive evaluation of P4DB in Section 7 using various OLTP benchmarks as well as an overview of related work in Section 8 and a summary in Section 9.

### 2 BACKGROUND

In the following, we discuss the relevant background on the typical programmable switches that are commercially available today.

#### 2.1 Programmable Switches

Programmable switches are an emerging trend on the market with specialized ASICs from vendors such as Intel [24] or Cavium [12]. Compared to traditional switches, programmable switches provide capabilities for flexible packet processing at line-rate of up to billion pkt/s. In-network processing (INP) enables advancements in the programmability of the data plane by providing a reconfigurable architecture [10] based on match-action tables. This increases the flexibility of switches to be used not only for data routing but also for offloading application logic to the switches by implementing customer match-action rules.

The de facto standard for programming the data plane [9] is P4 (Programming Protocol-Independent Packet Processors). P4 is a high-level language that allows to write match-action rules to express the processing of packets in the data plane [9]. Initially designed for programmable network switches, P4 is now used in a variety of systems which can also be used to process packets including SmartNICs or FPGAs [6, 38, 63].

P4 programs in general operate on packet headers and specify how to rewrite those headers. Although it uses a C-like syntax, it does not allow the use of pointers and has many other restrictions regarding floating-point numbers or loops to allow the processing of P4 programs at line-rate. These constraints can be overcome by common techniques like loop-unrolling, fixed-point arithmetic or dictionary encoding for strings. Switch vendors can also add different dedicated hardware-accelerators (e.g.: a checksum-engine or FPU) to the ASIC, which are then available through P4-externs.

#### 2.2 Protocol Independent Switch Architecture

For executing P4 programs, the commercially available programmable switches typically implement the so-called Protocol Independent Switch Architecture (PISA) as shown in Figure 2. PISA provides protocol-independence by allowing programmers to specify how a packet should be parsed and processed in a declarative manner using match-action tables as mentioned before. In P4DB, for example, we use custom match-action rules to implement the logic of the transaction engine on the switch. To process such application-specific match-action rules, PISAbased switches apply the following procedure: First, the header of a network packet is parsed in the ingress stage. These packet headers can include typical network information that is used for routing the packet along with metadata that can trigger more application-specific match-action rules that are executed in socalled match-action-units (MAUs) of PISA-based switches. In a PISA-based switch, multiple of these MAUs are typically linked together to process network packets in a pipeline-parallel manner with one packet per MAU stage. An important aspect is that match-action rules, if they depend on each other, must be placed in successive MAU stages. At the end of a switch pipeline, a packet is deparsed and sent over the network to its recipient.

# 2.3 Tofino Native Architecture

In this paper, we use a Tofino-based switch which is one of the commercially most successful programmable switches today. Based architecturally on PISA, it provides several extensions as part of the so-called Tofino Native Architecture (TNA) and programmable fixed-function components. For example, switch ports can be configured as loopback ports to cycle the same packet multiple times through the switch. This can be used to implement loops and thus execute data plane programs that can not be executed in a single pipeline pass. Stateful operations are also an interesting extension available in TNA [11, 32]. While these features are Tofino-specific, similar functions are provided by other vendors as well.

The ability of executing stateful operations in switches is a crucial part in P4DB, since it allows us to do transaction processing in a switch. The idea of these stateful operations in general is that they allow match-action rules to access state in SRAM through so-called register arrays. However, the total amount of memory, in today's switches is limited to a few megabytes only. While this seems to be a quite limiting factor, such an amount of memory is already enough to hold data for a few thousands of hot tuples, which has proven to be sufficient to significantly increase throughput of applications utilizing programmable switches [28, 29, 41, 68]. For example, the switch-programs of P4DB can store approximately 820K 8Byte hot tuples per pipeline. Moreover, there already exist other switch implementations based on FPGAs and high-bandwidth-memory that further alleviate this limitation [22].

Besides the limited switch memory, there exist other constraints on how data can be accessed from within P4 programs that are important aspects we need to consider when designing P4DB. For example, a first and important constraint, is that the register arrays (i.e., the SRAM) are partitioned over MAU stages and thus can only be accessed when the packet is processed by that stage. Moreover, to ensure the processing of packets at line-rate, multiple accesses (e.g., multiple read and write operations) from one packet to the very same register are not allowed [68]. Finally, a last but important limitation of stateful operations is that the order of register accesses in a P4 program needs to follow the placement order over the MAU stages in a pipeline (i.e., a data item stored in a register in an earlier MAU stage must be accessed first in the P4 program).

## **3 SYSTEM OVERVIEW**

In the following, we give an overview of P4DB. As shown in Figure 3, in an offline preparation step (Figure 3, left-hand-side) hot tuples

are offloaded from database nodes to the switch. Once offloaded, at runtime (Figure 3, right-hand-side) database nodes of the host DBMS then trigger transactions on the switch.

# 3.1 Offloading Data

A key challenge when offloading hot tuples to a programmable switch is to determine the data layout for the hot tuples on the switch. Storing tuples in registers of PISA-like switches such as the Tofino and supporting stateful operations on top comes with several constraints as discussed before. More precisely, hot tuples that should be stored on a switch have to be (1) assigned to registers in the order of how they are being accessed by transactions, and (2) a switch register that stores a tuple can only be accessed once per transaction and pipeline pass. In case these constraints can be satisfied by the data layout, transactions can be executed on the switch in a single pass through the pipeline. In Section 4 we hence present a new storage scheme for PISA-like switches called the declustered storage model, that aims to find a data layout for a given set of hot tuples and a set of transactions to maximize the number of transactions that can be executed in a single pass. Otherwise, if this is not the case, they require multiple passes through the switch pipeline, which negatively impacts the overall performance of executing switch transactions, as we discuss next.

In P4DB, we currently use a static approach to offload data; i.e., we decide based on access statistics of a representative workload which tuples are hot and thus should be offloaded to the switch. For the offline detection of hot tuples, we replay the transactions in a workload statement-by-statement to identify frequently accessed tuples. To identify those hot tuples at runtime, the partitioning manager of P4DB, which resides on each database node, keeps an index with the tuple identifiers (i.e., their primary keys) as well as other metadata. More details how this index is used at runtime and how secondary index lookups of hot tuples are supported is discussed in Section 6. Clearly, one could also use a more dynamic approach which monitors access frequencies at runtime to support potential shifts in the workload. However, this is an orthogonal aspect for future work.

#### 3.2 Processing Transactions

Once hot tuples are offloaded to a switch and the data layout (i.e., a mapping of tuples to registers) is defined, database transactions can be executed on the switch. In the following, we assume the case that database transactions can be partitioned into hot/cold transactions that either involve only hot tuples (on the switch) or cold tuples (on the host database nodes). Transactions that span over cold and hot tuples (called warm transactions) will be discussed in Section 6.

Hot Transactions. Hot transactions can be fully executed as inswitch transactions without coordinating with database nodes. The transaction logic of hot transactions is implemented as P4-programs in P4DB and deployed together with the data in the offloading step. For executing a hot (in-switch) transaction at runtime, database nodes of the host DBMS need to send a network message to the switch. This message includes all relevant transaction parameters such as transaction type that should be executed along with its input parameters. Once a hot transaction is executed, the results are sent back to the calling database node over the network.



Figure 3: P4DB accelerates distributed OLTP by exposing a programmable switch as an additional database node. To make use of the switch, P4DB first offloads hot tuples to the switch in an offline step. At runtime, transactions can then be executed at line-rate on these hot items on the switch.

The novelty of P4DB lies in the fact how transactions are being processed on the switch. The details this scheme will be discussed in Section 5. In a nutshell, our scheme enables that hot transactions (in most cases) can be executed using a single pass of the transaction through a switch pipeline. An important aspect is that single-pass transactions can be executed at line-rate and do not require any coordination on the switch while still guaranteeing serializability. In case transactions can not be executed in a single pass (which can be the case for more complex transactions with many reads and writes), P4DB falls back to a scheme for those transactions that requires multiple circulations through the switch pipeline and relies on coordination (as we discuss later). Hence, a goal of our storage scheme that we discuss next is to lay out the data on the switch to minimize the need for multi-pass transactions.

*Cold Transactions.* Different from hot transactions, cold transactions are executed completely by database nodes of the host DBMS without involving the switch. For processing transactions on cold tuples, not stored on the switch, P4DB currently uses a standard 2PC protocol for distributed transactions (with different deadlock detection modes) and a pessimistic lock based concurrency scheme to coordinate execution on each node. However, it is important to note that any other concurrency scheme could be used to execute cold transactions in the host DBMS (cf. Section 6).

# 4 DECLUSTERED STORAGE MODEL

In this section, we present our new declustered storage model for PISA-based switches like the Tofino switch. The goal of the declustered storage model is to execute hot transactions in a single pass through a switch pipeline.

#### 4.1 Goals of Data Layout

In the following, we first provide a high-level overview of how hot transactions are being executed in a switch. Afterwards, we derive the design of the data layout to enable single-pass transactions on the switch (i.e., which tuples are assigned to registers of the different MAU stages in a switch).

As shown in Figure 4, in P4DB each network packet in a switch pipeline represents a separate transaction. This assumption is reasonable, since (1) switches can execute billion packets per second, and (2) batching is often undesired in latency-critical OLTP. As a result, in each MAU stage of a switch pipeline, there is only a single transaction per cycle. Moreover, during processing, at each clock cycle transactions move to the next (subsequent) stage in a pipeline which is dictated by the switch execution model of PISA-based switches.



Figure 4: Pipelined processing of switch transactions. Packets represent transactions which advance in the pipeline once per clock cycle. In each cycle, a transaction can only access a single tuple of the current stage with one operation.

A first important implication that this execution model has on the data layout is that we only have one transaction per MAU stage with exclusive access to the registers of this MAU stage. This simplifies concurrency control for transactions on the switch, as we discuss next in Section 5. The first important constraint that is dictated by the switch memory model though is that a transaction can execute only one operation (read or write) to a register index which holds a single tuple. A stage usually can hold a few register arrays. Consequently, since transactions typically involve read/write operations on multiple tuples, a goal of the switch data layout is to allocate the tuples accessed by the same transaction into distinct register arrays of MAU stages. If this is not possible, a transaction needs to pass through the switch pipeline multiple times to execute all its required (read/write) operations.

Second, as mentioned before in Section 2, the access order of transactions has an influence on the assignment of tuples to registers of MAU stages. For example, if in Figure 4 *Transaction B* reads a tuple *A* and writes a tuple *B* using B = B + A, then *A* has to be located in a MAU stage before tuple *B*. Consequently, another goal of the switch data layout is to align the tuples stored in the register arrays of MAU stages with the access order imposed by transactions. Again, if it is not possible to reflect the order of accesses by a transaction in the data layout, a hot transaction needs to be executed using multiple passes through the switch pipeline.

Finally, there are hot transactions which always require multiple passes independent of the data layout (i.e., a single-pass execution is not possible at all). This is the case, when a transaction needs to execute multiple operations on the same tuple (e.g., a read and later a write of the same tuple), since only one operation per stage to the same tuple is allowed by the memory order of the switch. However, transactions that require multiple accesses to the same tuple are rather an exceptional case, as we see in our evaluation. Furthermore, as we will see in our experiments, P4DB can then still achieve significant benefits if such a case occurs compared to a traditional DBMS architecture without a switch.

# 4.2 Data Layout Problem

The goal of the declustered storage model is (as discussed before) to enable that hot transactions can be executed in a single pass by assigning tuples to switch registers that is optimal for the switch workload. More precisely, given a set of hot tuples and a set of hot transactions that should be executed on the switch, the goal of the declustered data layout is to maximize the number of hot transactions in the workload that can be executed in a single pass. Interestingly, this problem is related to declustered data access for disk-based databases that aims to spread data access over different disks [43–45].



Figure 5: Placement of tuples into switch stages and registers using the data layout algorithm on a graph representation. In a first step (left-hand-side), a max-cut is applied to the graph. Then, in a second step (right-hand-side), the partitions are then assigned to registers of different stages.

The problem of declustering data accesses is typically formalized as a graph problem. The general idea is to represent all tuples for which we aim to define a declustered data layout as nodes in the graph. If two tuples, represented by nodes U and V, are accessed together by one transaction, a new edge e(U,V) is added to the graph with a corresponding weight that describes their relative access frequency. In P4DB, we extend this model by using directed edges instead of edges without direction in order to model dependent operations that impose an order of access and thus an order of how to assign tuples to registers. More precisely, if an ordering dependency of operations (e.g., a read-write dependency) between two tuples represented by nodes U and V exist in one transaction, the edge is directed from U to V. If no such dependency exists, then a bidirectional edge is used between U and V to indicate that the tuples have no dependencies in the access order.

#### 4.3 Data Layout Algorithm

The data layout algorithm of P4DB uses this graph to assign hot tuples to registers of different MAU stages, with the goal to maximize the number of single-pass transactions. Recall from before, an edge e(U,V) with a high weight in the graph means that two tuples represented by the nodes are frequently accessed within the same transaction. Therefore, they should be placed into register arrays of separate MAU stages. In the following, we first ignore the edge directions and include them later in the discussion.

Overall, to find an optimal assignment of hot tuples to registers, we can apply a max-cut to partition nodes into N disjoint subgraphs that we aim to assign to N register arrays of the different MAU stages in a switch pipeline. The resulting partitions represent the set of tuples that should be allocated to the same register array of a stage. However, the max-cut problem of a graph is NP-complete [33]. Fortunately, many sophisticated approximation algorithms for the max-cut problem exist [18–20, 34, 54, 59]. For solving the max-cut problem of finding a partitioning of the hot-set on the switch in P4DB, we use MQLib [18] with additional constraints on the maximum size per partition, such that all tuples in a partition fit into the registers of a stage. Figure 5 (left-hand-side) shows an example of a graph that was constructed for six tuples (*T1* to *T6*) with weights representing access frequencies and the resulting max-cut for this graph into 3 different partitions.

After applying the max-cut, in the next step of the data layout algorithm, we need to assign each resulting partition to a register of the switch stage. For doing this, we now take directions of edges in the cut into consideration. In case the edges in a cut are either bidirectional edges or if all unidirectional edges in the cut between two partitions point into the same direction, then we assign the partitions with the outgoing edges into an earlier switch stage. For



Figure 6: Packet format of switch transactions. Each switch transaction corresponds to one network packet with a header and a variable amount of instructions, each of which defines an operation of a transaction.

example, in Figure 5 (right-hand-side) this is the case for the two cuts that separate the Partition 1 (blue) from Partition 2 (green) and Partition 2 (green) from Partition 3 (purple). In this case, we can simply do a topological ordering of the partitions based on the direction of the edges in the cuts and assign the partitions in the resulting order to switch stages.

In case that the edges in a cut point into different directions, we apply a slightly modified procedure. In this case, we first remove those edges in a cut such that all point in one direction and the sum of their weights is smaller than for the edges in the same cut that point in the other direction. The main idea is that these edges represent the accesses for which we have to violate the access order, resulting in a multi-pass transaction to support this order. Hence, we remove those edges with the total lower access frequency. The edges that remain in the cut between two partitions now all point to the same direction, and thus we can apply the procedure as discussed before to assign partitions to switch registers.

# 5 IN-SWITCH TRANSACTION PROCESSING

In this section, we describe how hot transactions are executed on the switch. We first explain the base case where a switch transaction can be executed in a single pass, before we then discuss transactions that might require multiple passes. Overall, the execution of hot transactions on the switch provides guarantees for *Atomicity*, *Isolation* and *Consistency* as we discuss next, while *Durability* is guaranteed by the integration into the host DBMS (cf. Section 6).

# 5.1 Single-Pass Transactions

In P4DB, as discussed in Section 4, we use a model where one transaction is mapped to one network packet. An interesting aspect of this model is that we get isolation and atomicity of transactions on the switch without additional coordination in case a transaction can be executed in a single pass.

First, we describe why our switch model actually guarantees isolation out-of-the-box. To better understand why this is the case, we refer back to Figure 4, which shows the pipelined execution of transactions in a switch pipeline. The important aspect is that our model dictates that there is only one transaction per MAU stage and transactions in different MAU stages are not re-ordered. Hence, the pipelined execution of transactions on the switch is equivalent to a serial execution order of transactions that results from the order as they are routed through the pipeline. For example, in Figure 4 the execution is equivalent to the serial order: *A*, *B*, *C*, and then *D*.

Second, another interesting aspect that results from the execution model in Figure 4 is that it also guarantees atomicity out-of-thebox, since no lock conflicts or deadlocks can occur on the switch. However, constraint checks to enable consistency could potentially still lead to aborts and thus would require additional efforts on the switch to roll back a transaction in case constraints are violated. To avoid these additional overheads for single-pass transactions, we use so-called constrained-writes of P4 to implement constraint checks. The main idea of a constrained-write in P4 is that a write is only executed if a predicate is satisfied. This allows us to support simple constraint checks at the end of a single-pass transaction on individual tuples, such as writing a transaction balance in Small-Bank only if its value is larger than zero at the end of the execution. For supporting more complex constraint checks though, one needs to fall back to a multi-pass execution scheme.

### 5.2 Multi-Pass Transactions

Hot transactions that are executed on a switch do not always satisfy the restrictions for single-pass execution. For example, Figure 6 shows a transaction that can not be executed in a single pass. As we can see, this transaction contains multiple operations, while the last two operations access the same register as the first two operations. Hence, this switch transaction needs to be executed in two passes. In the first pass, the three instructions at the beginning are executed and in the second pass the batch of the two remaining instructions. Yet, if the coordination-free execution would be used for multi-pass transactions, the isolation of transactions is no longer guaranteed.

To support isolation for multi-pass transactions, we thus provide a lock-based execution scheme in P4DB. In the following, we explain the naïve (fallback) approach that can always be used for multipass transactions and discuss further optimizations subsequently. The main idea of processing multi-pass transactions in the naïve scheme is that we introduce a so-called pipeline-lock to prevent the concurrent execution of multiple transactions in the same pipeline. The pipeline-lock is located in the first MAU stage of a pipeline. When a new transaction arrives at the switch pipeline and another multi-pass transaction is currently in the pipeline, the lock prevents its execution. This transaction is then scheduled for another pass through the switch using recirculation of the network packet.

Figure 7 shows an example with four transactions, where transaction *B* is a multi-pass transaction and all other transactions are single-pass. The pipeline-state is shown for multiple cycles horizontally to show the flow of transactions over time. In cycle 1, the single-pass transaction *A* is already in stage 1 when the multipass transaction *B* arrives in stage 0. Hence, transaction *A* is unaffected by the pipeline-lock that transaction *B* acquires. However, the pipeline-lock of transaction *B* then blocks the execution of following transactions *C* and *D* which are thus recirculated. Once transaction *B* starts its second pass though, it can directly unlock the pipeline-lock and thus *C* and *D* could be directly admitted to the pipeline. This is possible because all transactions follow serial execution and can not overtake transaction *B* in the pipeline.

Again, this multi-pass scheme guarantees isolation and atomicity of transactions since the pipeline-lock enables a serial execution and prevents violation of conflicts. Moreover, it also handles constraint checks to achieve consistency. To execute constraint checks in this model, additional passes of a transaction are required (once the



Figure 7: Processing of multi-pass transactions using a pipeline-lock. Each transaction checks the pipeline-lock in the first stage and can be recirculated to guarantee consistency of stored data, while one transaction is doing multiple passes through the switch pipeline.

execution of the normal operations is complete) to not only check the constraints, but eventually also to roll back changes in case a constraint is actually violated.

#### 5.3 Optimizations for Multi-Pass Transactions

This section gives a short overview of optimizations we implemented for multi-pass transactions in P4DB.

*Fine-grained Locking.* The switch's transaction throughput can be limited by the pipeline-lock that is used to preserve the consistency during the execution of multi-pass transactions. One optimization is clearly to use more fine-grained locking (e.g., one lock per MAU stage) with multiple locks for different MAU stages. With such a locking scheme, multiple transactions could run concurrently when they need locks for different disjunct registers in MAU stages. However, while a single pipeline-lock can be implemented in P4 by using a single atomic stateful operation on a register, it is not trivial to implement multiple lock instances for one pipeline, because a packet can not access a register in a stage twice or undo lock-acquisition for all locks if for example one lock failed.

An obvious solution is to utilize a bit representation of the lock, e.g. 64 bits, where each bit represents a separate lock. Unfortunately, Tofino's stateful register operations do not support bit-operations with variable masks as conditions for a compare-and-swap to test and set the lock. Moreover, using a complex if-cascade to check each bit separately is also not supported since this would result in code that can not be executed in a single cycle on the switch (i.e., the P4-compiler fails to compile such programs). However, with the current Tofino generation a 2-bit lock can be supported on the switch as shown in the pseudocode in Listing 1.

*Fast Recirculating.* Programmable switches allow recirculating packets by virtually connecting the output of a port to the input of the same port. This recirculation port is used as a queue for multi-pass (or single-pass) transactions which need to wait for their execution. A main challenge of P4DB is to minimize the time a pipeline-lock is held.

Therefore, our design uses two recirculation ports. The first port is used exclusively for transactions which own the lock on the switch. The second recirculation port is used for transactions which wait for execution. Since each port contains a separate queue for

```
struct lock_t { bit<8> left; bit<8> right; } // pair of locks
Register<...>(...) switch_lock;
RegisterAction<...>(switch_lock) try_lock = {
    void apply(inout lock_t value, out bit<1> rv) {
        if ((hdr.locks.left + value.left) == 2) {
            rv = 0; // fail, left lock is already locked
        } else if ((hdr.locks.right + value.right) == 2) {
            rv = 0; // fail, right lock is already locked
        } else {
            value.left = value.left + hdr.locks.left;
            value.right = value.right + hdr.locks.right;
            rv = 1; // requested locks acquired successfully
        }
    };
```

Listing 1: A 2-bit lock using a single register, in P4.

packets, transactions in the first recirculation port that hold a lock have a lower waiting time in the queue than the other (waiting) transactions, which reduces the time a lock is held by a transaction.

Finally, if a recirculation port receives more packets than it can handle, packets would be dropped. To prevent this, we actually split waiting transactions round-robin over multiple ports. To avoid starvation of recirculating transactions, each packet contains a recirculation counter, which can be used to prioritize the execution of transactions by the switch flow control as discussed before.

# 5.4 Processing Details

In this section, we finally discuss several important details of how (single- and multi-pass) transactions are being processed on the switch using network packets. A network packet needs to contain all information that is necessary for executing the transaction on the switch. More precisely, switch packets in P4DB contain a header which holds important processing information followed by a variable number of instructions representing the operations of a transaction as shown in Figure 6. These instructions are the operations that a transaction needs to run on the register arrays. Instructions can invoke trivial operations like read, writes or fixed point integer arithmetic as well as constrained-writes.

In the following, we now discuss the fields which hold processing information individually (shown in gray in Figure 6). The first field *is\_multipass* is a boolean flag which defines whether the incoming transaction can be executed in a single pass or not. Following that, a *locks* array shows the transaction engine which pipeline-locks need to be acquired in the first pass and freed in the last pass for multi-pass transactions. For single-pass transactions, this field does not indicate which pipeline-locks need to be acquired but rather which pipeline-locks need to be free such that the transaction can be admitted to a pipeline. The last field in the header is a counter *nb\_recircs* to track the number of times a packet is recirculated. It is incremented every time a transaction could not be executed and thus is recirculated. The switch flow control then uses this counter to prioritize long waiting transactions.

Finally, it is important to note that the processing information in packets must be initialized by database nodes that issue a hot transaction. Therefore, we keep all relevant information directly on the database nodes. For example, the information about the data layout (which tuple is stored in which MAU stage) is kept in an index structure redundantly per database node. That way, a database node can decide whether a transaction can be executed in a single pass or if multiple passes are required. More details about this are explained next in Section 6.



Figure 8: Processing of Warm Transactions. The switch transaction for the hot-set can only be sent when the operations on the cold tuples can not cause aborts anymore.
6 INTEGRATION WITH HOST DBMS

In this section, we discuss what is needed to integrate the switch execution into the host DBMS of P4DB.

# 6.1 Hot Transactions

As a first aspect, we discuss how the execution of hot transactions can be integrated into the host DBMS. Warm transactions that span across cold and hot tuples will be discussed next.

Execution Scheme. Most importantly, for executing hot transactions database nodes of the host DBMS need to trigger the execution of hot transactions. Hence, as a first extension, database nodes need to check for each transaction whether it accesses only tuples in the hot-set, otherwise the transactions is cold (or warm as we discuss next). To efficiently find out if a transaction accesses only hot tuples, we use an index that stores the keys of hot tuples. This index is replicated to all database nodes. Since it is accessed frequently and its size is small, it resides in CPU caches most of the time. Also, in the same index, we keep information in which stage and register a hot tuple is located on the switch. This information is used for a hot transaction to determine if it is a single-pass transaction or not, and thus to set the header fields of the network packet accordingly. Finally, secondary indexes are supported by keeping them on the database nodes in P4DB; e.g., for a lookup the secondary key is first mapped to a primary key to find out if a hot tuple is accessed. Moreover, the secondary indexes are updated based on the results of the switch transaction (which is possible in our case since switch transactions cannot fail).

Durability and Recovery. A second important aspect is that the database nodes handle the durability and recovery for switch transactions. To enable durability of switch transactions, database nodes use their local write-ahead log and append the operations for switch transactions they trigger to this log. In addition, to enable a correct recovery of the switch state from the different local logs of database nodes (in case of a switch failure), the switch adds a unique transaction ID to each switch transaction that it executes, which represents the (serial) execution order of the transactions on the switch. This ID is sent back together with the results of the read-and write-operations of a switch transaction to the database node in the response packet. The information is then appended to the write-ahead log of the database node.

In case the switch fails and needs to be recovered, the information in the local logs, which includes the globally ordered transaction IDs of the switch transactions, can then be used to recover a consistent state of the switch. A special case to consider for recovery are switch transactions that are in-flight (i.e., they are sent out but the result was not received by nodes). In P4DB, it is important to note that switch transactions count as committed before they are sent out since they cannot abort anymore. Therefore, a switch transaction and its intended read-/write-operations are appended to the log



Figure 9: Durability and Recovery of Warm Transactions. In this scenario, two warm transactions (T1 and T2) are executed. During execution, Node1 and the Switch both crash.

before the switch transaction is sent. The only info missing on database nodes at this point in the log is the unique transaction ID of switch transactions as well as the results of the read/write-operations. To handle this case, P4DB aims to restore the order from dependencies in the read/write-set (as we discuss in more detail for warm transactions). If no such dependency is detected, any order of switch transaction can be used during recovery for log replay.

#### 6.2 Warm Transactions

A second important aspect is how warm transactions are supported in P4DB that can span across both cold and hot tuples.

*Execution Scheme.* The main idea of the execution scheme for warm transactions is shown in shown in Figure 8. An important aspect of this scheme is the host DBMS first ensures that all operations on cold items will not abort before triggering the operations on the hot items on the switch using sub-transactions.

In P4DB, we use the following procedure to guarantee that operations on cold items cannot abort: In the first step, P4DB acquires the locks on all cold items on the database nodes. Once all locks on cold items are acquired, P4DB then executes operations on the cold items and checks all constraints using a sub-transaction that runs on the database nodes only. Once the sub-transaction on the cold items is ready to commit, P4DB sends out a network packet to the switch to trigger a sub-transaction, which then executes the operations on the hot items. After receiving the executed switch transaction, the updates on the hosts on cold items are committed as we discuss below (see 2PC).

It is important to note that the execution model above does not support cases where warm transactions need to access cold tuples after warm tuples (e.g., a write to a cold tuple that depends on a read of a hot tuple). We support this case in P4DB by additionally offloading those cold tuples to the switch during the offload phase of P4DB. For the cold tuples stored on the switch, we then use the very same processing scheme as for hot tuples (i.e., they are treated as if they were hot).

Durability and Recovery. Durability and recovery for warm transactions is handled similarly as for hot transactions. The only difference is that we need to consider additional failure cases. In the following, we discuss the most complex case where both a node and the switch that are involved in a warm transaction fail after switch transactions were sent out and before the result is received. Figure 9 shows an example for this case. In this scenario, the globally-unique transaction ID of T1 is missing in the log of Node1 and therefore the execution order of T1 and T2 cannot be restored during recovery, because it is not clear if T1 was executed before or after T2. However, we can still reconstruct the order of the two switch transactions for recovery by analyzing the logs of all database nodes for dependencies using the read/write-sets. For example, as shown in Figure 9, the order (T1 before T2) could be restored from the



Figure 10: Coordination of the distributed commit of warm transactions. Once the coordinator knows that the nodes (i.e., participants) can commit, it sends out the switch-transaction. After execution, the switch sends the commit decision via a multicast to the nodes.

read/write-set of T2 in the log of Node2 which indicates that T2 reads x=3 that must have been written by T1. If there were no dependency between T1 and T2, any order could be used <sup>1</sup>.

Integration with 2PC. Finally, a last important aspect that is different for warm transactions compared to hot transactions is that we need to coordinate the commit across multiple nodes (in case the cold part is distributed) and the switch. To coordinate the distributed commit of warm transactions, we extend the classical Two-Phase-Commit (2PC) protocol in P4DB as shown in Figure 10. Similar to the traditional 2PC, in our scheme a coordinator decides whether a warm transaction can commit based on the votes of all participating nodes similar to the normal 2PC. However, in addition if all nodes voted on commit, the switch sub-transaction is sent out by the coordinating node to the switch. The switch then executes the switch transaction and broadcasts the results using a multicast operation to all database nodes, which then finally commit their results without requiring additional round-trips. In case only one database node is involved in a warm transaction, we do not rely on 2PC. Instead, the switch sub-transaction can be sent out by the coordinator without executing a voting phase.

An interesting point of our integration with the 2PC protocol is that sub-transactions on the switch can directly commit once they are executed; i.e., sub-transactions on the switch do not need to wait for the commit on the database nodes. Note that this is not a problem since switch sub-transactions of warm transactions will commit in the same order as the host sub-transactions of warm transactions. The reason is that in case there is a conflict on colditems between two warm transactions, only the warm transaction which holds the locks on the conflicting cold items will send out the sub-transaction on the hot items. Otherwise, if there is no conflict of warm transactions on the cold tuples, then the sub-transactions on the hot-items can be executed and committed independently without coordination.

#### 7 EXPERIMENTAL EVALUATION

In the following, we present the results of our evaluation of P4DB. We first compare P4DB to other baselines using three OLTP benchmarks (YCSB, SmallBank, TPC-C). Afterwards, we then present microbenchmarks that show a more in-depth analysis of P4DB.

### 7.1 Setup and Baselines

*Setup.* The experimental setup consists of an 8-node cluster, where 6 of the nodes are equipped with two Intel(R) Xeon(R) Gold

<sup>&</sup>lt;sup>1</sup>To discuss additional failure cases, we plan to publish an extended technical report.

5120 2.2GHz CPUs and 2 nodes are equipped with two Intel(R) Xeon(R) Gold 5220 2.2GHz CPUs. Each node is connected to an Intel Tofino switch (BF2556X-1T) [5] via a 10G NIC (Intel X540-AT2). Hyper-threading is enabled for all nodes, but only one socket with 256GB of main-memory is used to avoid cross-socket communication to the NIC. The operating system is Ubuntu 18.04.1 LTS with Linux kernel 4.15.0 installed on all nodes. P4DB is implemented in C++20 and compiled with gcc-10. We utilize DPDK 20.08 [17] for optimized network performance and communication between the database nodes. The switch's control-plane logic of P4DB is implemented in C++ and the switch's data plane logic of P4DB is implemented in P4 and compiled using Intel SDE-9.3.0 [23].

P4DB Implementation. Our prototype implementation of P4DB<sup>2</sup> includes the switch that we integrated into a shared-nothing distributed in-memory DBMS as discussed in Section 6. To allow the execution of cold transactions in P4DB, we implemented two commonly used variants of Two-Phase Locking (2PL) with deadlock prevention: *NO\_WAIT* aborts a transaction as soon as a lock request is denied. For *WAIT\_DIE*, transactions are assigned a unique timestamp at transaction start and transactions only wait for a lock if it is owned by a transaction with a younger timestamp, otherwise it aborts. Similar to other recent papers [2, 21, 71], we do not include other deadlock detection algorithms such as *DL\_DETECT* since these schemes are more complex and typically do not provide any significant benefits.

*Baselines.* As a first baseline, we use the same prototype, however, our switch acts only as a traditional network switch and thus does not execute in-switch transactions (and hence the baseline is called *No-Switch*). As a second baseline, we used another variant of our prototype and use the switch only as a lock manager for hot tuples as suggested in [68] (called *LM-Switch*).

# 7.2 Workloads

For our experiments, we used three different OLTP benchmarks (YCSB, SmallBank, TPC-C). As the performance metric, we report throughput of committed (hot and cold) transactions per second while we also discuss latency at the end of the experiments. In the following, we describe each workload and how we defined the hot-set of tuples based on prior papers. Moreover, for all benchmarks, we vary the transactional load (i.e., the number of worker threads per database node) to analyze the effect of increased contention on the hot-set. In particular, we use configurations from 8 worker threads (normal load) to 20 worker threads (high load) per node, each executing a single transaction at a time.

*YCSB.* The Yahoo! Cloud Serving Benchmark [13] consists of a single table that is partitioned round-robin across all nodes. Similar to other papers [66, 72], we define a transaction as a group of 8 read/write-operations. We also populate the table with 1 billion entries consisting of an 8 byte key and 8 byte value.

In our evaluation, we use workloads A-C which are characterized by different ratios of read/write operations within a transaction to mimic update-heavy or read-heavy workloads (A - Update heavy 50/50, B - Read heavy 95/5, C - Read only 100/0 ). To simulate skew in the access patterns of tuples, similar to [67], we use a hot-set consisting of 50 key-value pairs per database node which receive 75% of all accesses.

*SmallBank*. The SmallBank benchmark models a banking application where transactions perform different operations on 1 or 2 customer accounts. The original workload contains 5 different types of transactions [1]. Similar to [16, 31], we additionally include a *Payment* transaction that transfers money between two accounts.

Compared to YCSB, SmallBank represents a workload with a fixed read-ratio of 15% and uses 1 million bank accounts. However, the difference to YCSB is that it contains read-dependent-writes (as well as some simple constraints), which makes it more complex to implement and shows the need for the declustered data layout in P4DB. To simulate skew in the access patterns of tuples, similar to [30], we used different hot-sets that vary in their sizes between 5 and 15 tuples per database node and that are accessed by 90% of all transactions.

*TPC-C.* The TPC-C benchmark [62] is the most complex benchmark in our evaluation and consists of 9 tables and 5 transaction types. Although TPC-C is designed to be highly partitionable by warehouse, it still contains several contention points. For example, the *NewOrder* transaction increments the global *next-order-id* per district and the *Payment* transaction updates the total balance per warehouse that causes contention on individual tables.

For our evaluation, we use a mix of only *NewOrder* and *Payment* transactions, similar to other papers [4, 66, 71], since these account for 90% of the transactional workload. Due to their complexity and the mix of cold and hot tuples within transactions, the logic of this workload requires warm transactions, unlike the previous benchmarks, and thus shows a different aspect of P4DB than the benchmarks before.

# 7.3 YCSB Experiments

In this experiment, we first validate our approach with YCSB as the most comprehensible OLTP benchmark. More precisely, we compare the throughput for YCSB using P4DB with its in-switch execution of hot transactions with the other two baselines; i.e., switch as lock-manager (*LM-Switch*) and a traditional distributed DBMS without using the switch for OLTP (*No-Switch*).

*Varying Contention.* In the first experiment, we compare the performance of P4DB and the baselines for a varying degree of contention. To increase the contention, we scale the number of worker threads as described before. Figure 11 (upper row) compares the speedup of P4DB w.r.t throughput over the *No-Switch* baseline, as well as the speedup of the *LM-Switch* over the same baseline. As we can see, P4DB outperforms both baselines and provides a speedup of up to 5× over the *No-Switch* baseline for the highest contention cases. The speedup of P4DB is the highest for workload A, due to its higher write-ratio causing decreased performance of both baselines.

Different from P4DB, the baseline which uses the switch as lockmanager (*LM-Switch*), does not yield any significant speed-up. In general, while *LM-Switch* enables high throughput for uniform workloads as shown in [68], this approach provides only minimal benefit for skewed workloads. The reason is that the overall latency of how long locks are held on contended items is not really reduced

<sup>&</sup>lt;sup>2</sup>https://github.com/DataManagementLab/p4db



Figure 11: YCSB Results — The speedup of P4DB and LM-Switch over No-Switch. P4DB yields the best improvements for increased contention (upper row) and an increased ratio of distributed transactions (lower row).



Figure 12: Benefits of switch transactions in YCSB. The switch executes the 75% hot transactions in the workload. In a traditional DBMS without switch a majority of those is aborted due to contention.

with this approach, which is very different from the coordinationfree model of P4DB.

Varying Distributed Transactions. Another aspect that is important in the context of P4DB is distributed transactions, since they could further amplify the effect of contention (due to higher access latencies of hot tuples). To investigate further the results from the previous experiment and to show the benefits of our in-switch approach for distributed transactions, we vary the degree of distributed transactions using the maximum number of worker threads from the previous experiment per node.

The results of this experiment can be seen in Figure 11 (lower row) which shows the speedup for P4DB and *LM-Switch* over *No-Switch*. Again, as before, P4DB provides a significant speedup while *LM-Switch* provides only limited gains for similar reasons. An interesting observation is that a higher degree of distributed transactions results in an increased throughput in P4DB leading to a speedup of more than  $6 \times$  for 100% distributed transactions due to the lower network latency (i.e., only ½ RTT) for hot tuples. Furthermore, at 0% distributed transactions, P4DB still provides a significant speedup over the *No-Switch* baseline. Here, the main reason is that the switch is much more efficient in processing transactions on the hot items than a normal database node.

*Break-down of Hot and Cold Tx's.* To better understand the effects of why P4DB dominates the throughput, we now present a throughput break-down of committed hot/cold transactions. For this experiment, we analyze the throughput (absolute results) for

20% distributed transactions and 20 worker threads per database node, as in the previous experiment. We decided to not show the *LM-Switch* baseline, since the break-down is very similar to the one of *No-Switch*.

Figure 12 shows the results of this experiment, where the fractions are marked for each bar. We see that for the update-heavy workload A, with *No-Switch* only a low fraction of hot transactions commits (dark grey, dark red) due to contention. P4DB, on the other hand, has a much higher commit ratio for hot transactions of around 75% (dark blue, dark yellow). For workload C, which is a pure read-only workload, the ratios for hot transactions in *No-Switch* and P4DB are the same (at approx. 75%) since no aborts happen for *No-Switch* on hot tuples. Nevertheless, P4DB's overall throughput is higher in all workloads, including workload C. Moreover, the switch throughput in P4DB is not affected by different read-write ratios of the workloads *A-C*, since all transactions can be executed (lock-free) in a single pass.

#### 7.4 SmallBank Experiments

In this experiment, we next compare the throughput for SmallBank using P4DB to *No-Switch*, which, as mentioned earlier, represents a more challenging workload as YCSB due to read-dependent-writes that could lead to multi-pass transactions. For executing this workload in P4DB, we use an optimal declustered data layout. We also show the effect of a non-optimal data layout in our microbenchmarks later. Furthermore, we dropped the *LM-switch* baseline from all subsequent experiments since it similarly shows only minimal gains over *No-Switch* as for YCSB.

*Varying Contention.* Similar as for YCSB, we first start with an experiment where we increase the load and thus the contention on a fixed set of hot tuples. Additionally, we use different hot-set sizes (5, 10, 15 hot tuples per database node) instead of different read/write ratios that we evaluated for YCSB.

The results of this experiment are shown in Figure 13 (upper row). In case for the smallest hot-set, P4DB again provides a speedup of up to 3× for the largest number of worker threads over the *No-Switch* baseline. Especially for the smallest hot-set, P4DB achieves high switch-throughput due to its single-pass execution model. Moreover, as the hot-set size increases and the contention decreases, P4DB still can provide a significant speedup of up to 2×.

*Varying Distributed Transactions.* Analogous to YCSB, we vary the amount of distributed transactions for SmallBank while fixing the number of workers to 20 per database node. Figure 13 (lower row) shows the results. With increasing fraction of remote transactions, we again see a speedup of P4DB over *No-Switch* even though the workload is more challenging as YCSB as discussed before due to read-dependent-writes. However, due to our declustered data layout of P4DB on the switch, all hot transactions in SmallBank can be executed efficiently in a single pass on the switch, resulting in a similar speedup as before for YCSB.

#### 7.5 TPC-C Experiments

In this experiment, we use TPC-C as a complex benchmark. As mentioned before, in TPC-C transactions cannot be implemented as solely hot or cold transactions, but it requires warm transactions that span across hot and cold tuples to include the switch



Figure 13: SmallBank Results – P4DB achieves a near-linear speedup for increased contention (upper row) and an increased ratio of distributed transactions (lower row).



Figure 14: TPC-C Results — Even though warm transactions are used, P4DB yields a significant speedup for increased contention (upper row) and an increased ratio of distributed transactions (lower row).

in the transaction processing. To be more precise, we offloaded all contended columns of the warehouse and district tables with write-accesses as well as stock columns of most ordered items to the switch. An interesting question thus is how much gains P4DB provides given this more challenging workload.

*Varying Contention.* The goal with TPC-C is to show whether using the switch as an accelerator when using warm transactions helps to increase overall system throughput. Again, we first scale the amount of worker threads per database node to show the effect of an increased contention while using a fixed ratio of 20% distributed transactions. We additionally vary the number of warehouses as they influence contention as well.

The results of this experiment are shown in Figure 14 (upper row). Overall, we see that with the highest contention configuration of 8 warehouses and 20 worker threads, P4DB still achieves a significant speedup of more than  $2\times$ . However, clearly, when increasing the number of warehouses to 16 and 32, the contention gets lower which decreases the gains. For example, with 32 warehouses, the speedup decreases to  $1.3\times$ .

Varying Distributed Transactions. This experiment shows the impact of a varying degree of distributed transactions when using warm transactions. For *NewOrder*, we varied the probability that an ordered item is located in a remote warehouse and for *Payment* that the paying customer is located at a remote warehouse.

Figure 14 (lower row) shows the results. Again, as for YCSB and SmallBank when increasing the ratio of distributed transactions, the access conflicts for *No-Switch* become more visible since more of the warm transactions need to wait for network round-trips and thus the latency for those transactions increases. In contrast, P4DB reduces the network round-trip cost for hot items, which leads to speedup increases of up to  $2.5 \times$  compared to *No-Switch*.

#### 7.6 Microbenchmarks

This section presents microbenchmarks to further understand details of the transaction execution on the switch when using P4DB. *7.6.1 Varying Hot/Cold Ratio.* In the previous experiments, we used a fixed ratio of hot to cold transactions. Thus in this microbenchmark, we further investigate the impact using different ratios. As a workload, we show the results for YCSB-A with 20% remote transactions, but the findings generalize to other workloads.

Figure 15a shows the absolute throughput, while Figure 15b shows the corresponding speedup of P4DB over *No-Switch*. As we see, when the fraction of hot transactions increases, the increased contention on hot tuples leads to a lower throughput of *No-Switch* caused by a higher abort rate. In stark contrast, P4DB's throughput increases with more hot transactions, resulting in a speedup of more than  $50 \times$  for 100% hot transactions. With 0% hot transactions, P4DB and *No-Switch* have a similar performance because the switch then only redirects cold transactions to database nodes and does not process any hot transactions.

7.6.2 Optimizations of Switch Processing. In this microbenchmark, we study the effect of the different optimizations we outlined in Section 5.3. For evaluating these effects in isolation, we use the hot (i.e., switch-only) transactions of the YCSB-A workload only. As a baseline (*Unoptimized*), we use a random data layout without any of the optimizations in Section 5.3 (i.e., fast recirculation and fine-grained locking turned off). We then turn on these optimizations one at a time, which optimize the efficiency of multi-pass transactions. Finally, we apply our optimal data layout, which aims to avoid multi-pass transactions as much as possible.

The results are shown in Figure 15c. As we can see, the first optimization called *Fast-Recirculate*, which creates a fast path for recirculating transactions that already own a pipeline-lock, leads to a speedup of approximately  $1.66 \times$  in P4DB. As a second optimization, we then applied fine-grained locking, which allows that two multi-pass transactions can be executed at the same time. This optimization together with the previous one results in a speedup of 2.6×. Finally, when applying our optimal data layout, we can further boost the performance since the need to use multi-pass transactions is significantly reduced. As a result, when turning on all optimizations and using the optimal data layout, we see an overall speedup of  $3.26 \times$  compared to the baseline (*No Optimization*).

*7.6.3 Single- vs. Multi-pass switch transactions.* As shown before, our data layout efficiently helps to improve the throughput of P4DB by enabling single-pass execution for more transactions. While in



(a) Skew: Throughput (b) Skew: Speedup (c) Optimizations Figure 15: Microbenchmarks — (a-b) shows the effect of a higher % of hot transactions and (c) shows the impact of optimizations for multi-pass transactions within P4DB.



Figure 16: Impact of the optimal data layout for the three workloads. SmallBank shows the highest benefits, while TPC-C is limited by warm transactions and hence the data layout has almost no effect.

the previous experiment we showed the effect only for the YCSB-A workload, in this experiment we use all three workloads to compare the effect of our data layout against a layout where tuples are randomly assigned to MAU stages (called worst case). We additionally report the average latency of transactions under the two data layouts. The results are shown in Figure 16.

We first discuss the effect of the data layout on the throughput. For SmallBank, we see that our data layout (called optimal case) has a much more significant effect on the throughput than for the YCSB-A benchmark. This effect is caused by a higher number of multi-pass transactions for a random data layout in SmallBank that cause an inferior performance. An optimal data layout, for TPC-C, has almost no effect compared to the random data layout because this workload uses warm transactions where the major limiting factor is the execution of cold sub-transactions rather than the execution of multi-pass transactions on the switch.

If we now look at the latency, we see similar effects. First, for YCSB and SmallBank, the optimal data layout enables that hot transactions can be executed in a single pass. Hence, the latency remains (almost) stable with increasing load (# of threads per server). This is in contrast to the random data layout, which causes multi-pass transactions (that require a pipeline-lock) which leads to increased latency under increased load. For TPC-C, latency increases for both data layouts instead since TPC-C uses warm transactions and here the latency is dominated by the cold sub-transactions. Overall, the average latencies are low (i.e., in the order of  $\mu$ s) for all benchmarks. Hot-set exceeds Switch's Capacity. An interesting question is 7.6.4 how P4DB performs if the hot-set size exceeds the switch's capacity. Hot-sets that exceed the switch capacity can be supported in P4DB by storing a part of the hot-set on normal nodes. For showing the effect of large hot-sets, we use the YCSB-A workload with increasing hot-set sizes. Furthermore, we use different tuple-width (i.e., values of different sizes) leading to a lower switch capacity in total number of rows the switch can store.



Figure 17: Throughput with growing hot-sets for YCSB-A. Different switch capacities result from different tuple widths. If the hot-set is beyond the switch capacity, the throughput in P4DB degrades gracefully and reaches the throughput of a DBMS without switch support.

Figure 17 shows how the throughput behaves while growing the hot-set above 4 different fixed switch capacities respectively, the largest being 650K rows. An important fact is that the throughput gracefully degrades when the hot-set outgrows the switch capacity in all cases (Note the log-scale of the x-axis). For example, for the largest switch capacity, P4DB can still achieve approximately 3 million txn/s for a hot-set of 2 million tuples which exceeds the switch capacity by more than  $2\times$ .

7.6.5 Latency-Breakdown for TPC-C. In Section 7.5, we showed that P4DB can achieve a performance gain of more than 2× speed-up for TPC-C under high contention (8 Warehouses). To understand where the benefits of using P4DB come from, we now show a break-down of the latency for committed transactions for a setup with 8 nodes each with 20 workers. The results in Figure 18a show the time spent in each component for our baseline without a switch and P4DB. Please note, that the total latency shown in this experiment is minimally increased due to profiling overhead compared to the other experiments before.

A first effect visible in Figure 18a is that P4DB can significantly reduce the time it takes transactions to synchronize on latches to check lock requests compared to the baseline (called *Lock Acquisition*). This is due to the execution scheme of P4DB which offloads contended items to the switch (i.e., warehouse, district or hot-stock items for TPC-C). Since P4DB can execute hot transactions without any locks, the average time spent for *Lock Acquisition* is reduced (i.e., locking is only needed for cold items). Furthermore, a second effect that leads to the benefits of P4DB is that *remote access* latency of transactions is reduced significantly since remote transactions on hot items only need half a round-trip to the switch instead of a full round-trip to other nodes. Overall, this results in less time needed to perform the actual work (i.e., remote reads/writes) for remote transactions, as shown in Figure 18a.

7.6.6 Existing Optimizations for Distributed Transactions and Contention. In this experiment, we analyze how existing optimizations for distributed transactions processing and contention compare to P4DB. To show the effects of existing optimizations, in this experiment we start with a *Plain 2PL* protocol using two-phase commit (2PC) for distributed transactions without any additional optimizations that clearly results in the lowest throughput. Afterwards, we add typical optimizations for distributed transactions incrementally to the baseline including P4DB. We use the TPC-C workload which is the most complex workload from the previous experiments. Due to space restrictions, we show the results for 8 warehouses only



(a) Latency-Breakdown (b) Existing Optimizations Figure 18: Further Microbenchmarks for TPC-C: (a) Shows the latency-breakdown to understand where the benefits of using P4DB come from. (b) shows effects of existing optimizations for processing distributed and hot transactions. (which results in the highest level of contention). For a higher number of warehouses we see similar results. The results can be seen in Figure 18b compared to P4DB.

As a first optimization, we apply optimal data partitioning to improve data locality of distributed transactions (called +Opt. Part. in Figure 18b). In order improve locality in OLTP various partitioning schemes such as [15] have been proposed. For showing the general effect of locality in this experiment, we use two different TPC-C configurations - one with 80% and another one which results in 20% remote transactions respectively. As a second optimization, we additionally apply a contention-centric execution scheme on top of the optimal data partitioning. For the experiment, we added a recent scheme called Chiller [71] to our baseline which builds on 2PL but it aims to increase throughput on hot and contented items by deploying a two-region execution scheme with early lock release for contended items instead of using plain 2PL. Alternatively to Chiller, we offload the hot-set of hot tuples to the switch and use P4DB for executing transactions on the contended items. The results show that the optimizations significantly improve the 2PL/2PC protocol but P4DB still outperforms these optimizations clearly.

#### 8 RELATED WORK

In the following, we discuss related work using INP for OLTP (which is the main target in this paper) in detail. In addition, other related work of using INP for OLAP exists such as [8, 39, 55, 56].

INP Support for OLTP. While INP has been used to support OLTP, to the best of our knowledge P4DB is the first approach to offload OLTP processing to programmable switches. NetCache [29] proposed a key-value store on a programmable switch to cache hot-items. Different from this approach, P4DB supports the execution of multiple operations as a transaction directly on the switch, while NetCache only supports simple (read/write) operations. In addition, NetCache uses a different approach for reads and writes, which leads to inferior performance of update-heavy workloads. Another direction is to support replication schemes by INP. For example, Harmonia [73] implements a replication protocol for OLTP DBMSs in a switch data-plane to overcome scalability issues. Moreover, NetChain [28] proposes a replication protocol across multiple programmable switches. All these approaches are clearly orthogonal to P4DB. Furthermore, another orthogonal work is [27] which proposes an INP approach called Transaction Triaging that manipulates streams of transactions by batching, re-ordering, steering, and protocol conversion of transactions on a programmable switch to improve efficiency. Finally, NetLock [68] proposes a central lockmanager which directly processes lock-requests of hot items in the

switch's data-plane. However, as we have shown in Section 7.3, NetLock provides inferior performance for OLTP workloads under high contention when compared to P4DB.

*OLTP under High Contention.* Another line of work aims to reduce contention with specialized protocols in the host DBMS. For example, one direction is to reduce contention through enforcing determinism to part of, or all of the concurrency control (CC) unit [14, 31, 60]. Another direction is Quro [65], which re-orders operations inside transactions in a centralized DBMS with 2PL to reduce lock duration of contended data. While this is relevant in the context of this paper, almost all these works deal with single-node DBMSs and do not have the notion of distributed transactions. In the context of distributed DBMS, there has recently been some work to leverage high-speed networks and data partitioning to reduce contention in OLTP [71]. While this work clearly targets a similar problem, it is orthogonal and thus an interesting direction would be to combine such an approach with P4DB in the future.

SmartNICs for OLTP. Previous work leveraged SmartNICs as coprocessors to offload distributed transaction processing [46, 58]. However, processing hot data on SmartNICs has several drawbacks compared to programmable switches. First, programmable switches are connected to all network nodes centrally and are not bandwidth-limited like SmartNICs with only a few network ports (e.g. 2x50GbE). Processing at full aggregated network-bandwidth is enabled by P4DB's pipelined and abort-free transaction execution engine implemented on the switch ASIC. Second, and more importantly in the OLTP context, programmable switches can be reached in only 1/2 network latency compared to remote SmartNICs which require a full roundtrip to a node which hosts the SmartNIC. While certain SmartNICs that use general-purpose processors for INP allow more complex per-packet processing, they thus also require coordination of data-accesses by multiple threads (e.g. through locking, similar to normal databases nodes) which presents a major bottleneck for hot items as we have shown. Furthermore, upcoming switch models provide increasing resources for computation and memory [25, 26] or even FPGAs as on-chip co-processors [3].

# 9 CONCLUSION

In this paper we presented P4DB, a first-of-its-kind approach that leverages a programmable switch to accelerate highly contented OLTP workloads by offloading transaction processing to the switch. For this, P4DB exposes the switch as an additional database node that not only stores hot tuples, but also implements pipelined and abort-free transaction processing directly in the network's data plane. The two major benefits of our in-switch design are (1) access in only ½ latency and (2) aggregated processing bandwidth of all connected nodes. In our evaluation, we have shown that P4DB can thus effectively mitigate contention, i.e., skew on hot tuples, yielding a substantial increase in throughput due to its pipelined transaction execution on a switch.

#### ACKNOWLEDGMENTS

This work was partially funded by the German Research Foundation (DFG) under the grants BI2011/1 & BI2011/2 (DFG priority program 2037) and the DFG Collaborative Research Center 1053 (MAKI). We thank Intel for their valuable technical support.

#### REFERENCES

- M. Alomari, M. Cahill, A. Fekete, and U. Rohm. The cost of serializability on platforms that use snapshot isolation. In 2008 IEEE 24th International Conference on Data Engineering, pages 576–585. IEEE, 2008.
- [2] R. Appuswamy, A. C. Anadiotis, D. Porobic, M. K. Iman, and A. Ailamaki. Analyzing the impact of system architecture on the scalability of oltp engines for high-contention workloads. *Proceedings of the VLDB Endowment*, 11(2):121–134, 2017.
- [3] APS2140D (Tofino+FPGA). https://www.aps-networks.com/products/aps2140d/, 2021.
- [4] T. Bang, N. May, I. Petrov, and C. Binnig. The tale of 1000 cores: An evaluation of concurrency control on real (ly) large multi-socket hardware. In *Proceedings* of the 16th International Workshop on Data Management on New Hardware, pages 1–9, 2020.
- Barefoot. Barefoot Tofino APS Networks BF2556X-1T-A1F. https://stordirect. com/shop/switches/25g-switches/aps-networks-bf2556x-1t-a1f/, 2019.
- [6] Barefoot Networks: The world's fastest most programmable networks. https://www.gdt.id.au/~gdt/presentations/2016-07-05-questnetsdn/papers/barefoot2016--barefoot--the-worlds-fastest-and-mostprogrammable-networks.pdf, 2021.
- [7] C. Binnig, A. Crotty, A. Galakatos, T. Kraska, and E. Zamanian. The end of slow networks: it's time for a redesign. *Proceedings of the VLDB Endowment*, 9(7):528-539, 2016.
- [8] M. Blöcher, T. Ziegler, C. Binnig, and P. Eugster. Boosting scalable data analytics with modern programmable networks. In *Proceedings of the 14th International* Workshop on Data Management on New Hardware, pages 1–3, 2018.
- [9] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, et al. P4: Programming protocol-independent packet processors. ACM SIGCOMM Computer Communication Review, 44(3):87–95, 2014.
- [10] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz. Forwarding metamorphosis: Fast programmable match-action processing in hardware for sdn. ACM SIGCOMM Computer Communication Review, 43(4):99–110, 2013.
- [11] M. Budiu and C. Dodd. The p416 programming language. ACM SIGOPS Operating Systems Review, 51(1):5–14, 2017.
- [12] Cavium-XPliant (R) family of programmable Ethernet switches. https://www. openswitch.net/cavium/, 2021.
- [13] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with ycsb. In *Proceedings of the 1st ACM symposium on Cloud computing*, pages 143–154, 2010.
- [14] J. Cowling and B. Liskov. Granola: low-overhead distributed transaction coordination. In 2012 {USENIX} Annual Technical Conference ({USENIX} {ATC} 12), pages 223–235, 2012.
- [15] C. Curino, E. P. C. Jones, Y. Zhang, and S. R. Madden. Schism: a workload-driven approach to database replication and partitioning. 2010.
- [16] D. E. Difallah, A. Pavlo, C. Curino, and P. Cudre-Mauroux. Oltp-bench: An extensible testbed for benchmarking relational databases. *Proceedings of the* VLDB Endowment, 7(4):277–288, 2013.
- [17] DPDK. https://www.dpdk.org/, 2021.
- [18] I. Dunning, S. Gupta, and J. Silberholz. What works best when? a systematic evaluation of heuristics for max-cut and qubo. *INFORMS Journal on Computing*, 30(3):608–624, 2018.
- [19] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified np-complete problems. In Proceedings of the sixth annual ACM symposium on Theory of computing, pages 47–63, 1974.
- [20] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- [21] R. Harding, D. Van Aken, A. Pavlo, and M. Stonebraker. An evaluation of distributed concurrency control. *Proceedings of the VLDB Endowment*, 10(5):553–564, 2017.
- [22] J. Hofmann, L. Thostrup, T. Ziegler, C. Binnig, and A. Koch. High-performance in-network data processing. In International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures, ADMS@ VLDB, 2019.
- [23] Intel P4 Studio. https://www.intel.com/content/www/us/en/secure/design/ confidential/products-and-solutions/networking-and-io/programmableethernet-switch-products/p4-suite/p4-studio.html, 2021.
- [24] Intel (R) Programmable Ethernet Switch Products. https://www.intel. com/content/www/us/en/products/network-io/programmable-ethernetswitch.html, 2021.
- [25] Intel (R) Tofino (TM) 1. https://www.intel.com/content/www/us/en/products/ network-io/programmable-ethernet-switch/tofino-series.html, 2021.
- [26] Intel (R) Tofino (TM) 2. https://www.intel.com/content/www/us/en/products/ network-io/programmable-ethernet-switch/tofino-2-series.html, 2021.

- [27] T. Jepsen, A. Lerner, F. Pedone, R. Soulé, and P. Cudré-Mauroux. In-network support for transaction triaging. *Proceedings of the VLDB Endowment*, 2021.
- [28] X. Jin, X. Li, H. Zhang, N. Foster, J. Lee, R. Soulé, C. Kim, and I. Stoica. Netchain: Scale-free sub-rtt coordination. In 15th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 18), pages 35–49, 2018.
- [29] X. Jin, X. Li, H. Zhang, R. Soulé, J. Lee, N. Foster, C. Kim, and I. Stoica. Netcache: Balancing key-value stores with fast in-network caching. In Proceedings of the 26th Symposium on Operating Systems Principles, pages 121–136, 2017.
- [30] A. Kalia, M. Kaminsky, and D. G. Andersen. Fasst: Fast, scalable and simple distributed transactions with two-sided ({RDMA}) datagram rpcs. In 12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16), pages 185–201, 2016.
- [31] R. Kallman, H. Kimura, J. Natkins, A. Pavlo, A. Rasin, S. Zdonik, E. P. Jones, S. Madden, M. Stonebraker, Y. Zhang, et al. H-store: a high-performance, distributed main memory transaction processing system. *Proceedings of the VLDB Endowment*, 1(2):1496–1499, 2008.
- [32] P. G. Kannan, R. Joshi, and M. C. Chan. Precise time-synchronization in the data-plane using programmable switching asics. In Proceedings of the 2019 ACM Symposium on SDN Research, pages 8-20, 2019.
- [33] R. M. Karp. Reducibility among combinatorial problems. In Complexity of computer computations, pages 85–103. Springer, 1972.
- [34] Y.-H. Kim, Y. Yoon, and Z. W. Geem. A comparison study of harmony search and genetic algorithm for the max-cut problem. Swarm and evolutionary computation, 44:130–135, 2019.
- [35] T. Kraska, M. Hentschel, G. Alonso, and D. Kossmann. Consistency rationing in the cloud: Pay only when it matters. *Proceedings of the VLDB Endowment*, 2(1):253–264, 2009.
- [36] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1):14–76, 2014.
- [37] V. Krishnaswamy et al. Relative serializability: An approach for relaxing the atomicity of transactions. J. Comput. Syst. Sci., 55(2):344-354, 1997.
- [38] R. Kundel, L. Nobach, J. Blendin, H.-J. Kolbe, G. Schyguda, V. Gurevich, B. Koldehofe, and R. Steinmetz. P4-bng: Central office network functions on programmable packet pipelines. In 2019 15th International Conference on Network and Service Management (CNSM), pages 1–9. IEEE, 2019.
- [39] A. Lerner, R. Hussein, P. Cudre-Mauroux, and U. eXascale Infolab. The case for network accelerated query processing. In CIDR, 2019.
- [40] E. Levy, H. F. Korth, and A. Silberschatz. A theory of relaxed atomicity. In Proceedings of the tenth annual ACM symposium on Principles of distributed computing, pages 95–110, 1991.
- [41] X. Li, R. Sethi, M. Kaminsky, D. G. Andersen, and M. J. Freedman. Be fast, cheap and in control with switchkv. In 13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16), pages 31–44, 2016.
- [42] Y. Li, I.-J. Liu, Y. Yuan, D. Chen, A. Schwing, and J. Huang. Accelerating distributed reinforcement learning with in-switch computing. In 2019 ACM/IEEE 46th Annual International Symposium on Computer Architecture (ISCA), pages 279–291. IEEE, 2019.
- [43] D.-R. Liu and S. Shekhar. A similarity graph-based approach to declustering problems and its application towards parallelizing grid files. In Proceedings of the eleventh international conference on data engineering, pages 373–381. IEEE, 1995.
- [44] D.-R. Liu and S. Shekhar. Partitioning similarity graphs: A framework for declustering problems. *Information Systems*, 21(6):475–496, 1996.
- [45] D.-R. Liu and M.-Y. Wu. A hypergraph based approach to declustering problems. Distributed and Parallel Databases, 10(3):269–288, 2001.
- [46] M. Liu, T. Cui, H. Schuh, A. Krishnamurthy, S. Peter, and K. Gupta. Offloading distributed applications onto smartnics using ipipe. In *Proceedings of the ACM* Special Interest Group on Data Communication, pages 318–333. 2019.
- [47] L. Michael, W. Nejdl, O. Papapetrou, and W. Siberski. Improving distributed join efficiency with extended bloom filter operations. In 21st International Conference on Advanced Information Networking and Applications (AINA'07), pages 187–194. IEEE, 2007.
- [48] A. Pavlo, C. Curino, and S. Zdonik. Skew-aware automatic database partitioning in shared-nothing, parallel oltp systems. In Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, pages 61–72, 2012.
- [49] A. Pavlo, E. P. C. Jones, and S. Zdonik. On predictive modeling for optimizing transaction execution in parallel oltp systems. *Proc. VLDB Endow.*, 5(2):85–96, Oct. 2011.
- [50] O. Polychroniou, R. Sen, and K. A. Ross. Track join: distributed joins with minimal network traffic. In *Proceedings of the 2014 ACM SIGMOD international conference* on Management of data, pages 1483–1494, 2014.
- [51] D. R. Ports and J. Nelson. When should the network be the computer? In Proceedings of the Workshop on Hot Topics in Operating Systems, pages 209–215, 2019.
- [52] A. Quamar, K. A. Kumar, and A. Deshpande. Sword: scalable workload-aware data placement for transactional workloads. In *Proceedings of the 16th international conference on extending database technology*, pages 430–441, 2013.

- [53] W. Rödiger, S. Idicula, A. Kemper, and T. Neumann. Flow-join: Adaptive skew handling for distributed joins over high-speed networks. In 2016 IEEE 32nd International Conference on Data Engineering (ICDE), pages 1194–1205. IEEE, 2016.
- [54] S. Sahni and T. Gonzalez. P-complete approximation problems. Journal of the ACM (JACM), 23(3):555–565, 1976.
- [55] A. Sapio, I. Abdelaziz, A. Aldilaijan, M. Canini, and P. Kalnis. In-network computation is a dumb idea whose time has come. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, pages 150–156, 2017.
- [56] A. Sapio, I. Abdelaziz, M. Canini, and P. Kalnis. Daiet: a system for data aggregation inside the network. In *Proceedings of the 2017 Symposium on Cloud Computing*, pages 626–626, 2017.
- [57] A. Sapio, M. Canini, C.-Y. Ho, J. Nelson, P. Kalnis, C. Kim, A. Krishnamurthy, M. Moshref, D. Ports, and P. Richtarik. Scaling distributed machine learning with in-network aggregation. In 18th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 21), pages 785–808, 2021.
- [58] H. N. Schuh, W. Liang, M. Liu, J. Nelson, and A. Krishnamurthy. Xenic: Smartnicaccelerated distributed transactions. 2021.
- [59] K. Seo, S. Hyun, and Y.-H. Kim. An edge-set representation based on a spanning tree for searching cut space. *IEEE Transactions on Evolutionary Computation*, 19(4):465-473, 2014.
- [60] A. Thomson, T. Diamond, S.-C. Weng, K. Ren, P. Shao, and D. J. Abadi. Calvin: fast distributed transactions for partitioned database systems. In *Proceedings of* the 2012 ACM SIGMOD International Conference on Management of Data, pages 1–12, 2012.
- [61] M. Tirmazi, R. Ben Basat, J. Gao, and M. Yu. Cheetah: Accelerating database queries with switch pruning. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 2407–2422, 2020.
- [62] TPC-C Specification. http://www.tpc.org/tpc\_documents\_current\_versions/pdf/ tpc-c\_v5.11.0.pdf, 2021.
- [63] H. Wang, R. Soulé, H. T. Dang, K. S. Lee, V. Shrivastav, N. Foster, and H. Weatherspoon. P4fpga: A rapid prototyping framework for p4. In Proceedings of the Symposium on SDN Research, pages 122–135, 2017.

- [64] X. Wei, J. Shi, Y. Chen, R. Chen, and H. Chen. Fast in-memory transaction processing using rdma and htm. In Proceedings of the 25th Symposium on Operating Systems Principles, pages 87–104, 2015.
- [65] C. Yan and A. Cheung. Leveraging lock contention to improve oltp application performance. Proceedings of the VLDB Endowment, 9(5):444–455, 2016.
- [66] X. Yu, G. Bezerra, A. Pavlo, S. Devadas, and M. Stonebraker. Staring into the abyss: An evaluation of concurrency control with one thousand cores. 2014.
- [67] X. Yu, Y. Xia, A. Pavlo, D. Sanchez, L. Rudolph, and S. Devadas. Sundial: harmonizing concurrency control and caching in a distributed oltp database management system. *Proceedings of the VLDB Endowment*, 11(10):1289–1302, 2018.
- [68] Z. Yu, Y. Zhang, V. Braverman, M. Chowdhury, and X. Jin. Netlock: Fast, centralized lock management using programmable switches. In Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication, pages 126–138, 2020.
- [69] E. Zamanian, C. Binnig, T. Harris, and T. Kraska. The end of a myth: Distributed transactions can scale. *Proc. VLDB Endow.*, 10(6):685–696, Feb. 2017.
- [70] E. Zamanian, C. Binnig, and A. Salama. Locality-aware partitioning in parallel database systems. In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, pages 17–30, 2015.
- [71] E. Zamanian, J. Shun, C. Binnig, and T. Kraska. Chiller: Contention-centric transaction execution and data partitioning for modern networks. In *Proceedings* of the 2020 ACM SIGMOD International Conference on Management of Data, pages 511–526, 2020.
- [72] E. Zamanian, X. Yu, M. Stonebraker, and T. Kraska. Rethinking database high availability with rdma networks. *Proceedings of the VLDB Endowment*, 12(11):1637– 1650, 2019.
- [73] H. Zhu, Z. Bai, J. Li, E. Michael, D. R. K. Ports, I. Stoica, and X. Jin. Harmonia: Near-linear scalability for replicated storage with in-network conflict detection. *Proc. VLDB Endow*, 13(3):376–389, Nov. 2019.
- [74] T. Ziegler, S. Tumkur Vani, C. Binnig, R. Fonseca, and T. Kraska. Designing distributed tree-based index structures for fast rdma-capable networks. In Proceedings of the 2019 International Conference on Management of Data, pages 741–758, 2019.