

AdaChain: A Learned Adaptive Blockchain

Chenyuan Wu
University of Pennsylvania
wucy@seas.upenn.edu

Bhavana Mehta
University of Pennsylvania
bhavanam@seas.upenn.edu

Mohammad Javad Amiri
University of Pennsylvania
mjamiri@seas.upenn.edu

Ryan Marcus
University of Pennsylvania
rcmarcus@seas.upenn.edu

Boon Thau Loo
University of Pennsylvania
boonloo@seas.upenn.edu

ABSTRACT

This paper presents *AdaChain*, a learning-based blockchain framework that adaptively chooses the best permissioned blockchain architecture in order to optimize effective throughput for dynamic transaction workloads. *AdaChain* addresses the challenge in the Blockchain-as-a-Service (BaaS) environments, where a large variety of possible smart contracts are deployed with different workload characteristics. *AdaChain* supports automatically adapting to an underlying, dynamically changing workload through the use of reinforcement learning. When a promising architecture is identified, *AdaChain* switches from the current architecture to the promising one at runtime in a way that respects correctness and security concerns. Experimentally, we show that *AdaChain* can converge quickly to optimal architectures under changing workloads, significantly outperform fixed architectures in terms of the number of successfully committed transactions, all while incurring low additional overhead.

1 INTRODUCTION

Blockchain systems, in particular, *permissioned* blockchain systems, have enabled a new class of data center applications, ranging from contact tracing [51], crowdworking [14], supply chain assurance [15, 59], and federated learning [52]. The popularity of these services has motivated cloud providers, e.g., Amazon [2, 3], IBM [8], Oracle [9], and Alibaba [63], to offer *Blockchains-as-a-Service* (BaaS) [25].

BaaS offerings have resulted in a large variety of possible smart contracts being deployed. Different smart contracts may exhibit different workload characteristics, such as read/write ratios, skewness of popular keys, compute intensity, etc. To address these variations in workloads, there has been a proliferation of permissioned blockchain systems, e.g., Tendermint [43], Fabric [16], Fabric++ [56], Fabric# [54], Streamchain [37], and ParBlockchain [11]. These blockchain systems present significant variation in architectural design, including the sequence in which ordering, execution and validation are done, the number of transactions in a block, stream processing (with no blocks), and the use of reordering and early aborts.

Past studies [22, 31] have shown that different blockchain architectures and hyperparameter settings are optimal for different workloads with varying properties (e.g. system load, write ratios, skewness, and compute intensity). We experimentally confirmed this observation: Figure 1 shows the performance of various blockchain

architectures across four different workloads, showcasing significant variations in throughput. For example, for Workload A¹, which requires high levels of computation, an XOv architecture with re-ordering (e.g., Fabric++ [56]) provides the best throughput of all the tested architectures. On the other hand, for Workload D, which requires significantly less computation but has higher skewness, an OXII architecture (e.g., ParBlockchain [11]) demonstrates the highest throughput. Clearly, there is some dependency between workload characteristics and the optimal blockchain architecture for that workload.

Currently, BaaS providers must choose a single architecture to offer customers, potentially resulting in poor performance, as no single architecture provides dominant throughput. Even when the user has control over the blockchain architecture, choosing the right architecture and parameters is not easy given the large configuration space. Moreover, in a BaaS setting, the workload may fluctuate and change, as different tenants scale up or down their smart contracts deployments, and client requests fluctuate with different patterns throughout the day. Of course, one could imagine building a static mapping from workload characteristics to optimal blockchain architectures – but this mapping would (1) be expensive to compute, (2) depend on the underlying hardware, (3) still be suboptimal for workloads that shift unexpectedly over time, and (4) require recomputing the mapping each time a new blockchain architecture is developed.

In this paper, we propose *AdaChain*, a reinforcement learning-based blockchain framework that chooses the best blockchain architecture and sets appropriate parameters in order to maximize effective throughput for dynamic transaction workloads. Experimentally, we show that *AdaChain* is not only able to select optimal or near-optimal configurations for a wide variety of workloads, but its reinforcement learning approach also allows it to quickly adapt to new hardware, new storage subsystems, and new unanticipated workload changes on the fly.

In order to build an adaptive blockchain, *AdaChain* relies on two key innovations. First, it models the selection of a blockchain architecture as a contextual multi-armed bandit problem, a well-studied reinforcement learning problem with asymptotically optimal results [39]. This formulation allows *AdaChain* to apply classical algorithms, such as Thompson sampling [23], to select blockchain architectures in a way that minimizes *regret* (the difference between the performance of the chosen architecture and the optimal architecture). *AdaChain* will strategically test different architectures to learn which ones are well-suited to the user’s workload. It learns which architectures work best by observing the characteristics of

¹Details about each workload can be found in Tables 2 and 3.

the workload and the effective throughput of the system. When the workload changes, AdaChain notices drops in throughput, and can automatically adjust the blockchain architecture and parameters to maximize performance, all without any user intervention.

Second, AdaChain introduces protocols to switch from one blockchain architecture to another in a live system, while maintaining strong serializability properties. This switching protocol is not only required for AdaChain to function (multi-armed bandits generally require making multiple decisions before the optimal is reached), but also enables a new class of blockchains that can more-or-less seamlessly transition between different architectures to support the shifting workloads in the real-world. Intuitively, the switching protocol works by splitting switching decisions between two paths. In the normal path, all nodes agree to switch to the same new architecture after a certain number of blocks have been committed, while in the slow path, all nodes switch to the same architecture after failing to make progress on processing transactions for a certain amount of time.

Specifically, this paper makes the following contributions.

- **Learned adaptive blockchain.** To the best of our knowledge, AdaChain is the first blockchain system to support *automatically adapting to an underlying, dynamic workload*. Through careful modeling of the states, actions, and objective function, AdaChain’s use of reinforcement learning makes it the first blockchain system to *learn from its mistakes* and self-correct.
- **Multi-architecture switching.** Additionally, we also present the first blockchain system capable of switching from one architecture to another at *runtime* in a way that respects correctness and security concerns.
- **Analysis of architecture impact on blockchain performance.** We perform a large-scale measurement examining the relationship between architecture choice and blockchain performance. We implemented a wide range of blockchain architectures, and through a suite of workload parameters, we identified architecture configurations and runtime settings that significantly impact performance improvements. Our experiments highlight the large state space, which renders manual heuristics difficult to achieve. The workloads, architectures and measured performance will be publicly available to aid future research.
- **Prototype and performance evaluation.** We have developed a prototype of AdaChain, which will be publicly available under an open-source license. Our evaluation results on CloudbLab demonstrate that AdaChain can converge quickly to optimal architectures under changing workloads, significantly outperform existing fixed architectures, and incur low additional overhead.

2 ARCHITECTURE LANDSCAPE

To motivate AdaChain, we first examine, both experimentally and intuitively, why different blockchain architectures work better for different workloads: in Section 2.1, we will highlight a number of blockchain architectures, and illustrate their advantages and disadvantages. The point here is not that some blockchain architectures are “better” or “worse” than others, but rather that each blockchain architecture performs well under some conditions and poorly under

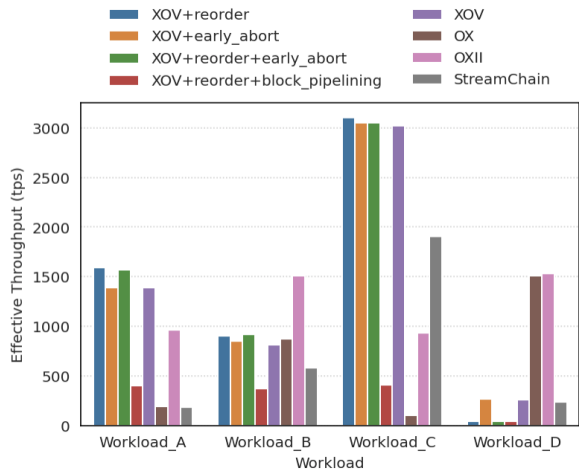


Figure 1: Performance of different blockchain architectures under various workloads. The performance of different architectures can vary significantly between workloads. Workloads and architectures are described in Table 1 and 2.

others. In Section 2.2, we argue that a blockchain that can adaptively switch between multiple architectures is able to achieve “best of all worlds” performance.

2.1 Blockchain Architectures and Workloads

Table 1 lists representative blockchain systems and their architectures, where the design space consists of seven performance optimizations (P1-P7) and two correctness dimensions (C1-C2). Figure 1 shows their corresponding performance under four different workloads. Here, we use *effective* throughput as the performance metric, which measures the number of successfully committed transactions per second.

The workloads A to D are characterized in Table 2. BaaS workloads embody a large extent of variations. For instance, different transactions might invoke different percentages of write operations to the underlying key-value store, as represented by the *write ratio*. These transactions might also contend to access or update the same set of popular keys (or hot keys), as indicated by the *contention level*. In addition, the runtime *load* on a BaaS can be determined by the frequency of issued transactions by each client and the number of active clients varying with time. Last, *compute intensity* is an important characterization of BaaS workloads, as pointed out by [34, 54, 62]. This is because permissioned blockchains support a wide range of applications, some of which are compute-intensive (e.g., those that provide security and correctness guarantees for machine learning applications).

Below, we briefly describe the design principles of each architecture and explain the intuition behind why the performance of each architecture can vary under different workloads.

Order-Execute (OX). The order-execute architecture has been widely used in permissioned blockchain systems such as Tendermint [43], Quorum [24], Chain Core [4], Multichain [35], Hyperledger Iroha [7], and Corda [5]. In the OX architecture, transactions are totally ordered and batched into blocks and then transactions

Architecture	Rep. System	P1. Block Size	P2. Early Exec.	P3. Dependency Graph	P4. Early Abort	P5. Cross-Block Conflicts	P6. Parallel Exec.	C1. MVCC	C2. Isolation
OX	Tendermint [43]	tunable	✗	✗	✗	-	✗	✗	strong serializable
OXII	ParBlockchain [11]	tunable	✗	✓	✗	-	partial	✗	strong serializable
XOV	Fabric [16]	tunable	✓	✗	✗	V	fully	✓	strong serializable
XOV++	Fabric++ [56]	tunable	✓	✓	X, O	V	fully	✓	strong serializable
XOV#	Fabric# [54]	tunable	✓	✓	O	O	fully	✗	serializable
XOV	StreamChain [37]	1	✓	✗	✗	V	fully	✓	strong serializable

Table 1: Comparing design principles of existing permissioned blockchain architectures. Here, P stands for performance, C stands for correctness, X stands for execution, O stands for ordering, and V stands for validation.

Workload	Write Ratio	Contention Level	Load	Compute Intensity
A	low	high	high	high
B	moderate	high	moderate	low
C	moderate	low	high	very high
D	high	very high	moderate	very low

Table 2: Characterizing workloads A, B, C and D. Specific workload parameters are presented in Table 3.

of a block are executed sequentially. As a result, the OX architecture does not require an MVCC validation phase, which is used to resolve conflicts between transactions, and hence, no transactions will be aborted due to conflicts. As shown in Figure 1, this design principle makes OX outstanding at workload D, where transactions are write-heavy and contentious, i.e., transactions update a small set of hot keys. On the other hand, OX performs comparatively poorly on workloads A and C, which are compute-intensive. Due to the lack of parallel execution mechanisms, OX cannot take advantage of the multi-core processing power of modern servers.

Order-Parallel Execute (OXII). In the OXII architecture, used by ParBlockchain [11], transactions are first totally ordered and batched into blocks. OXII then constructs a dependency graph for transactions within a block based on their positions. Specifically, if t_i is ordered before t_j , and the pair of transactions have either WR, RW, or WW conflicts, OXII constructs a directed edge from t_i to t_j . This dependency graph is then used in the execution phase to execute certain transactions in parallel, i.e., a transaction can be executed once all its predecessors have finished execution. Given a higher level of execution parallelism than OX, OXII performs better than OX under computation-heavy workloads such as A and C. Note that even for a given workload, OXII requires careful tuning of block size; a large block results in high overhead in dependency graph construction, while a small block results in less parallelism and higher communication overhead.

Execute-Order-validate (XOV). Hyperledger Fabric [16] presents the XOV architecture (which was first introduced by Eve [38] in the context of Byzantine fault-tolerant SMR) by switching the order of the ordering and execution phases such that transactions are simulated fully in parallel before being ordered in the ordering phase. Since it utilizes early execution, XOV requires an MVCC validation phase to invalidate all transactions that are simulated on stale data, and commits only the validated transactions to the world-state and the blockchain ledger. This early execution enables XOV to perform well on contention-free workloads such as C. On the other hand, XOV demonstrates poor performance under contentious and write-heavy workloads, such as B and D, due to the high percentage of invalidated transactions.

XOV with early abort and reordering (XOV++). The XOV++ architecture, as introduced in Fabric++ [56], follows the XOV paradigm but with some modifications. First, a dependency graph is constructed in the ordering phase to capture RW conflicts between each pair of transactions within the same block. When the graph is constructed, all elementary cycles in the graph are aborted greedily. Unlike OXII, which utilizes the graph for concurrency control, XOV++ uses the graph for transaction reordering; when there is a RW conflict between t_i and t_j , it (re)orders t_i before t_j in the block. Second, it adopts early abort techniques in both the simulation and ordering phases. Whenever XOV++ detects that a transaction operates on stale data, XOV++ immediately aborts that transaction without waiting for the final MVCC validation. As an effect of transaction reordering, XOV++ has outstanding performance on workload A, where the conflicts are reconcilable given a low write ratio. On the other hand, XOV++ performs poorly on workload D with a near-zero effective throughput. This is because, under a contentious and update-heavy workload, very few conflicts can be reconciled through reordering. Moreover, reordering becomes more expensive when there are a large number of cycles in the dependency graph, resulting in more pending blocks and, thus, more transactions that simulate on stale data.

XOV with serializable isolation (XOV#). The XOV# architecture, presented in Fabric# [54], is mainly different from XOV and XOV++ in that XOV# is serializable, while XOV and XOV++ are strong serializable. To achieve this isolation level, XOV# incrementally constructs a dependency graph that keeps track of all dependencies, including those that span across blocks in the ordering phase. Once a transaction is ordered, XOV# immediately drops this transaction if there is a dependency cycle involved. The resulting acyclic schedule is guaranteed to be serializable, thus, no extra MVCC validation is needed in XOV#. To ensure a fair comparison with other architectures, we run XOV# under the strong serializability isolation level while keeping the remaining design dimensions the same as the original XOV# (the XOV+reorder+block_pipelining bar). XOV# performs worse than vanilla XOV in all workloads A to D due to the overhead of maintaining a large dependency graph and detecting cycles. This suggests that the performance improvement reported in Fabric# is mainly due to a more relaxed isolation level.

Stream XOV. StreamChain [37] switches from block processing to stream transaction processing. Specifically, StreamChain follows the XOV paradigm while fixing the block size to 1. The motivation behind stream processing is simple: while the original, permissionless blockchains were forced to use proof of work (PoW) consensus techniques to maintain fault tolerance, a permissioned blockchain

environment allows more efficient consensus protocols to be used. Thus, stream processing can reduce transaction latency. In terms of effective throughput, StreamChain has relatively good performance when the workload is lightweight or not contentious, such as in workloads B and C. Otherwise, the high block construction overhead in terms of cryptographic operations and excessive disk I/Os leads to a large number of pending blocks in StreamChain, making incoming transactions simulate on stale data.

StreamChain also highlights that the *parameters* of a given architecture can impact performance [22]. A large block size leads to higher block formation overhead and latency, while a small block size results in higher communication and disk overhead.

2.2 The Case for Adaptivity & Learning

The previous section can be summarized as follows: depending on workload and hardware characteristics, the performance of a given blockchain architecture can vary drastically. We thus argue that *there is no one-size-fits-all architecture*. From Figure 1, we observe that not only is there no single dominant architecture, but some architectures that perform well on one workload can end up performing quite poorly on another.

Adaptivity. Even if one precomputed the optimal blockchain configuration for a particular workload, one would still have to deal with: (1) shifts in workload characteristics could render the precomputed decisions suboptimal, (2) changes in the underlying hardware may cause a different architecture to become optimal, and (3) the development of new blockchain architectures would require re-computing a new optimal. Moreover, BaaS providers must ensure acceptable performance even when (1) multiple customers have distinct workloads, (2) those workloads vary throughout the day, and (3) the performance of hardware may vary. A blockchain that can switch between architectures can thus adapt to the wide variety of changes that occur in modern BaaS environments.

Learning. One may consider the design of simple heuristics to map workloads to optimal blockchain configurations. For example, a simple rule like “when compute intensity is high, use XOV, otherwise use OXII” would indeed improve upon a single workload. However, a heuristic that achieved optimal or near-optimal performance for a wide variety of workloads would be much harder to craft. Even if one invested the time to craft such a heuristic, the heuristic would quickly become outdated, as changes in hardware and new blockchain architectures emerge.

Instead of undertaking the Sisyphean task of designing such a heuristic, we instead propose using machine learning techniques to *learn* the heuristic on-the-fly. By using reinforcement learning techniques, we can quickly search an ever-expanding set of blockchain configurations for any workload on any hardware. Such a system can *learn from its mistakes* and progressively improve.

3 ADACHAIN OVERVIEW

At a high level, AdaChain contains two key components: a machine learning model (the *learning agent*) which guides AdaChain towards better and better blockchain architectures, and an *architecture switching* mechanism that allows AdaChain to near-seamlessly transition from one blockchain architecture to another while ensuring correctness and security.

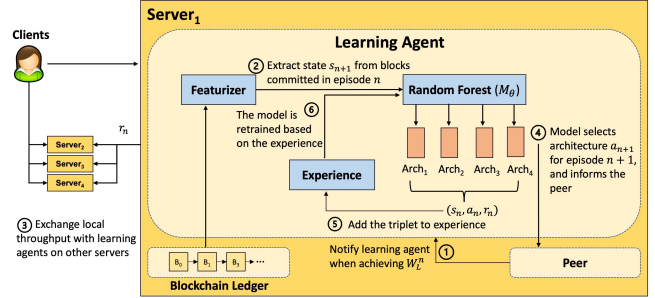


Figure 2: Architecture of AdaChain. For readability, we only present the internals of server₁.

Learning agent. AdaChain’s learning agent models the problem of selecting a blockchain architecture as a contextual multi-armed bandit (CMAB) problem [64]: periodically, AdaChain examines the most recent properties of the workload (*context*), and then selects one of many blockchain architectures (*arms*). After making the selection, it observes the effective throughput of the newly selected architecture (*reward*). To be successful, AdaChain must balance the *exploration* of new, untested architectures with *exploiting* past experience to maximize throughput – without a careful balance of exploration and exploitation, AdaChain risks failing to discover an optimal configuration (too much exploitation), or performing no better than random (too much exploration). We select this CMAB formulation (as opposed to generalized reinforcement learning models) because CMABs are exceptionally well-studied, and many asymptotically-optimal algorithms exist to solve them [10, 23].

Details about the learning agent are provided in Section 4.

Switching architecture. AdaChain utilizes a switching protocol that allows it to switch from one blockchain architecture to another in a distributed fashion across all nodes in the blockchain deployment, while transactions are ongoing. AdaChain achieves this by splitting switching decisions between two paths, a normal path in which all nodes agree to switch to the same new architecture after a certain number of blocks have been committed, and a slow path in which all nodes switch to the same new architecture after failing to make progress for a certain amount of time.

Details about AdaChain’s switching protocol are in Section 5.

AdaChain workflow overview. Figure 2 shows the overall architecture of AdaChain. Similar to other permissioned blockchains, AdaChain consists of a set of distributed servers, where each server runs a peer process. The peers are responsible for transaction processing and appending the constructed blocks to the blockchain ledger. In addition to the peer process, AdaChain introduces a separate learning agent process on each server. The learning agent is responsible for finding the current optimal architecture according to the workload, and guiding its local peer to switch to that optimal architecture.

AdaChain operates in *episodes*, where within one episode, the blockchain architecture remains unchanged. When the learning agent finds an architecture candidate, it instructs the peer to use that architecture for the next episode. Each episode is marked by the completion of a constant number of transactions ($\Delta N_{\text{episode}}$), including invalidated transactions. This episode design ensures

AdaChain does not stuck in a bad architecture for a long time even when the fraction of invalidated transactions is high due to conflicts.

Below, we describe how the learning agent proposes the architecture for episode $n + 1$ in detailed steps. Although our discussion below focuses on the internals of server₁, the same procedure happens simultaneously on every blockchain server.

Step 1: Notifying the learning agent. In episode n , the peer notifies its local learning agent when the number of committed blocks has reached a certain watermark. The notification also includes the local performance measurement r_n in episode n .

Step 2: Featurization. Since learning agents are distributed across different servers, the state (i.e., some features that capture the workload) that they need in order to make a decision should also be distributed. In AdaChain, states are not only distributed, but also decentralized as no single entity controls the state. This is possible with negligible overhead due to two key insights: (1) the blockchain ledger contains rich information about the workload, and thus is a good source of raw data for featurization; (2) the ledger is naturally decentralized and consistent across peers. Thus, once the agent is notified by the peer, its featurizer extracts the state s_{n+1} from blocks committed in episode n . Features are described in detail in Section 4.2.

Step 3: Exchanging performance measurements. The locally observed performance is different across different peers, and malicious peers could even manipulate the local measurement. To ensure each honest server has the same architecture for episode $n + 1$, the learning agent on server₁ exchanges the local measurement r_n with learning agents on every other server, so as to agree on performance measurement.

Step 4: Estimating the performance for each architecture. The predictive model M_θ predicts the performance of each architecture candidate under state s_{n+1} , and selects architecture a_{n+1} that is predicted to have the best performance. The learning agent then informs the peer to switch to a_{n+1} for episode $n + 1$.

Step 5: Building experience buffer. Once an r_n is obtained, the learning agent adds the (s_n, a_n, r_n) triplet to the experience buffer. Note that s_n and a_n are derived prior to the start of episode n .

Step 6: Retraining. The predictive model M_θ is periodically re-trained based on the experience buffer, creating a feedback loop. As a result, AdaChain’s predictive model improves, and AdaChain more reliably picks the best architecture for the observed state.

Assumptions. In short, AdaChain can adapt itself according to the workload and hardware setup to continually improve performance. Moreover, AdaChain is an *online* learned system that does not require a separate and cumbersome data collection process prior to deployment. Our current design makes two assumptions. First, in AdaChain, similar to many other permissioned blockchain systems [13, 24, 43], each node serves as both the ordering and execution (endorser) node. This, however, is in contrast to Hyperledger Fabric and its variants that separate endorsing and ordering roles. Second, AdaChain is designed for a homogeneous setup, where different servers have access to similar resources. While having these two assumptions in place simplifies the system design, they have been used in real-world BaaS deployments. Removing these assumptions is an avenue for future work.

4 LEARNING ALGORITHMS

In this section, we discuss AdaChain’s learning approach in detail. We first formalize AdaChain’s learning problem as a contextual multi-armed bandit problem, and then discuss our selected algorithm, Thompson sampling, for solving such problems. We next describe the predictive model used by AdaChain, followed by the specific state and action space design.

Contextual multi-armed bandits (CMABs). In a contextual multi-armed bandit problem, an agent periodically makes decisions in a number of episodes, enumerated by n . In each episode, the agent selects an action a_n based on a provided state s_n , and then receives a reward r_n . The agent’s goal is to select actions in a way that minimizes *regret*, i.e., the difference between the reward from the chosen action and the reward from the optimal action. CMABs assume that each episode is independent² from each other, and that the optimal decision depends only on the state s_n . As described in Section 3 and 5, in order to be responsive to workload changes, there are no pending blocks across different episodes in AdaChain. Thus, each episode in AdaChain can also be considered to be independent (although caching effects may bring a small amount of dependence between episodes).

AdaChain’s formulation. AdaChain uses *effective* throughput as the performance metric P to maximize, which is the number of successfully committed transactions per second. For each episode, it must select an architecture to use. AdaChain’s goal is to select the best architecture (in terms of effective throughput) in the family of available architectures A , given the current perceived workload $w \in W$. We call this selection function $S : W \rightarrow A$. We formalize the goal as a regret minimization problem, where the regret r_n for an episode n is defined as the difference between the effective throughput of the architecture selected by AdaChain and the ideally optimal architecture as presented in equation 1.

$$r_n = \max_{a \in A} P(w, a) - P(w, S(w)) \quad (1)$$

We use effective throughput as the performance metric since it is the dominant metric used by previous studies. However, the optimization goal above can be easily generalized to other performance metrics, such as latency or even user-defined service level objectives. We leave the study and evaluation of such performance metrics to future work.

Thompson sampling. While there are many algorithms to solve contextual multi-armed bandit problems, we select Thompson sampling for its simplicity: at the start of each episode, we train a model based on our current experience, and then select the best action as predicted by the model. In order to train the model, Thompson sampling deviates from traditional ML techniques. Normally, models are trained by finding the model parameters that are most likely given the training data (i.e., maximum likelihood estimation). This works well when the training data is a representative sample of the population. Unfortunately, in the context of a multi-armed bandit, the experience set is not a representative sample: it contains only data for actions we have previously selected. As a result, if we trained our model using standard maximum likelihood techniques, our agent would assume that our experience was representative,

²Contextual multi-armed bandit algorithms have been shown to be effective even when these condition do not strictly hold [23].

and would only *exploit* past knowledge and would perform very little *exploration* (i.e., testing actions that the current data suggests may be sub-optimal). Alternatively, if we wanted our agent only to explore, we could set the model weights to random values, ensuring a random prediction at each episode.

How can we balance these competing goals? We want to *exploit* the information we have gathered in the past, but we want to avoid getting stuck in local maxima by *exploring* new possibilities as well. The beauty of Thompson sampling is that an optimal balance can be struck between exploration and exploitation by slightly modifying the training procedure: instead of selecting the model parameters that are most likely given the data, instead we *sample model parameters proportionally to their likelihood given the training data*. More formally, we can define maximum likelihood estimation as finding the model parameters θ that maximize likelihood given experience E : $\arg \max_{\theta} P(\theta | E)$ (assuming a uniform prior). Instead of maximizing likelihood, Thompson sampling simply samples from the distribution $P(\theta | E)$. This means that if we have a lot of data suggesting that our model weights should be in a certain part of the parameter space, our sampled parameters are likely to be in the part of the space. Conversely, if we have only a small amount of data suggesting that our model weights should be in a certain part of the parameter space, we may or may not be sample parameters in that part of the space during any given episode.

Of course, the “test loss” of a sampled θ may be higher than the “test loss” of the most likely parameters: this is a key difference between supervised and reinforcement learning. Concisely, sampling from $P(\theta | E)$ instead of finding the most likely parameters $\arg \max_{\theta} P(\theta | E)$ means that the likelihood of us choosing particular model parameters is proportional to the quantity of evidence for those model parameters in our experience.

Comparison with supervised learning. The contextual bandit problem formulation is critical for AdaChain to be practical: AdaChain is an *online* system that learns from its mistakes, without requiring a separate training data collection process prior to deployment. If we simply wanted to select the architecture with the best-expected performance, one naïve way would be training a predictive model in a standard supervised fashion. However, as our value model might be wrong, we might not always pick the optimal architecture, and, as we would never try alternative strategies, we would never learn when we are wrong. In other words, we could become trapped in a local maxima.

Training a value model in a standard supervised fashion also requires a time-consuming data collection process because an accurate value model requires *complete* data: there cannot be large “holes” in either the state space or action space when enumerating them to gather data. Thus, all the potential workloads, blockchain architectures, and hardware setups must be known *a priori*. This is not realistic in a BaaS because: (1) unexpected changes in hardware could happen when the cloud provider introduces new CPUs, larger RAM, or transitioning from local disk to remote storage engine, etc., (2) users might exhibit unexpected workload behaviors, and (3) the action space keeps growing as new blockchain architectures emerge. When these changes happen, the previously trained value model will be wrong and misguide the peers, unless the time-consuming data collection process is repeated.

4.1 Predictive Model

AdaChain uses random forests [18] as the predictive model due to their good performance on data sets of moderate sizes and fast inference. The model takes the state (i.e., workload) concatenated with action (i.e., architecture choice) as input, and outputs the predicted performance.³ Thus, given a state, AdaChain enumerates the action space and uses the model to predict the performance of each action. AdaChain then chooses the action with the best-predicted performance to be carried out. Once there is a tie on the best-predicted performance, AdaChain breaks the tie randomly to avoid local maxima.

Integrating random forests with Thompson sampling requires the ability to sample model parameters from $P(\theta | E)$. The simplest technique (which has been shown to work well in practice [49]) is to train the model as usual, but only on a bootstrap [17] of the training data. In other words, the random forest is trained using $|E|$ random samples drawn with replacement from experience E , inducing the desired sampling properties. AdaChain uses this bootstrapping technique for its simplicity.

In AdaChain, each node’s learning agent starts with the same random seed when AdaChain is launched. Thus, since the state of a certain episode is the same across peers (as mentioned in Section 3 and 5), with the predictive model’s deterministic training and inference, each honest agent chooses the same blockchain architecture in the same episode.

4.2 State Space

In AdaChain, the state represents properties of the client workload. AdaChain captures the state space using the four features below. To ensure the accuracy of feature extraction, all aborted or invalidated transactions are still written to the ledger with a validity flag (similar to [36]). Below, we assume a window of blocks b_i to b_j in the ledger are read by the learning agent for featuring the current state.

Write ratio. We observe that counting the write ratio in terms of write accesses to the key-value store is not effective for predicting performance. Thus, AdaChain measures the write ratio at the transaction level: once a transaction writes to the key-value store, it is viewed as a write transaction. The write ratio is the ratio of write transactions to the number of all transactions during b_i and b_j .

Hot key ratio. AdaChain measures the frequency that each key is accessed during b_i and b_j . It then takes the frequency corresponding to the hottest key to be the hot key ratio.

Transaction arrival rate. AdaChain timestamps each transaction upon its first arrival to the system. AdaChain first measures the number of all transactions from b_i to b_j , denoted as N , and then derives the transaction arrival rate using $\frac{N}{ts_j - ts_i}$, where ts_i represents the arrival timestamp of the first transaction in b_i and ts_j represents that of the last transaction in b_j .

Execution delay. AdaChain uses average execution delay of all transactions in the period of b_i to b_j .

4.3 Action Space

In AdaChain, the action space consists of different blockchain architectures. One naïve approach to represent the action space is

³This corresponds to a *value based model*. A *policy model*, in which the predictive model predicts simultaneously the probability of each action being optimal, might be an interesting direction for future work.

to give each architecture a one-hot encoding. However, from the random forest’s perspective, this approach makes two semantically close architectures totally unrelated, resulting in ineffective splits, and thus poor prediction accuracy. For example, assume XOV is represented by vector $(1, 0, \dots, 0)$ in the one-hot encoding. Random forest might split on the first dimension in the vector, i.e., XOV is its left child, while everything non-XOV is its right child. Each child’s performance will be predicted using the average performance of that child. Clearly, XOV++ and StreamChain might have a relatively close performance to XOV, but they will always fall into a wrong child node and their predicted performances are wrongly averaged.

Thus, AdaChain chooses to first featurize the blockchain architectures to maintain the semantic information of their design. Feature engineering an optimal representation of blockchain architectures is a difficult and inexact task. Instead of attempting to design an all-encompassing representation that captures every dimension of blockchain architectures, we instead selected a simple representation based on our intuition of the most important properties. We leave investigating alternative representations to future work.

AdaChain therefore captures the action space using three main features: block size, early (speculative) execution, and dependency graph construction. Block size is a scalar variable, representing the number of transactions within a block. The block size in AdaChain is also equal to the batch size in the consensus protocol. To limit the growth of action space, the block size can not exceed 1,000, which is larger than typical block sizes used in blockchain systems, and we further discretize the block size using paces. Early execution and dependency graph construction are both binary variables. Thus, AdaChain’s action space consists of 100 choices in total.

We do not consider parallel execution as a feature because it can be derived from the two previous features (i.e., early execution and dependency graph construction): (1) early execution of transactions happens fully in parallel; (2) the goal of constructing a dependency graph is to execute independent transactions in parallel.

5 SWITCHING ARCHITECTURES

This section discusses the architecture switching mechanism of AdaChain. We first introduce the normal path of operations, followed by our timeout-based mechanism in the slow path. Lastly, we outline why the switching mechanism is resilient to attacks.

5.1 Normal Path

Algorithm 1 presents the normal path of operations. Each server in AdaChain runs Algorithm 1 in a distributed fashion in order to carry out architecture switching. Here, S is the set of blockchain servers, i stands for the index of the server, n stands for the current episode, and $\Delta N_{\text{episode}}$ and ΔN_{learn} are two constant hyperparameters. At a high level, the normal path introduces two watermarks: a low watermark (W_L^n) that triggers the learning phase, and a high watermark (W_H^n) that marks the end of an episode.

The untrustworthiness of participants in a blockchain system prevents us from relying on a centralized entity to featurize the state and measure the reward. Thus, inspired by the PBFT protocol [20], AdaChain conducts them in a decentralized fashion. Upon reaching the low watermark W_L^n , each server $i \in S$ records its locally observed throughput p_i^n of episode n and featurizes the state

Algorithm 1 Normal path

```

> On each server  $i$ 
1: Upon index of local last committed block  $b_{\text{last}}$  reaching  $W_L^n$ 
2:   Record performance  $p_i^n$ 
3:   Extract features  $f_i^{n+1} = (w_i^{n+1}, c_i^{n+1}, r_i^{n+1}, e_i^{n+1})$  from block  $W_H^{n-1}$  to  $W_L^n$ 
4:   Multicast  $\langle \text{CHECKPOINT}, n, i, e_i^{n+1}, p_i^n \rangle_{\sigma_i}$  to all servers
> On the leader server  $l$ 
5: Upon receiving valid CHECKPOINT messages from a quorum  $Q$  of  $2f + 1$  servers
6:   Compute  $e^{n+1} \leftarrow \text{median}\{e_j^{n+1} \mid j \in Q\}$ 
7:   Compute  $p^n \leftarrow \text{median}\{p_j^n \mid j \in Q\}$ 
8:   Multicast  $\langle \langle \text{PROPOSE}, e^{n+1}, p^n \rangle_{\sigma_l}, C \rangle$  to all servers
> On each server  $i$ 
9: Upon receiving a PROPOSE message from the leader
10: if  $e^{n+1}$  and  $p^n$  are valid (based on  $C$ ) then
11:   Multicast  $\langle \text{ACCEPT}, n, i, e^{n+1}, p^n \rangle_{\sigma_i}$  to all servers
12: Upon receiving valid matching ACCEPT messages from  $2f + 1$  different servers
13:   Multicast  $\langle \text{COMMIT}, n, i, e^{n+1}, p^n \rangle_{\sigma_i}$  to all servers
14: Upon receiving valid matching COMMIT messages from  $2f + 1$  different servers
15:   Add  $p^n$  to experience and derive action  $a_{n+1}$  based on  $f^{n+1}$ 
16: if  $T^n$  transactions have been committed then
17:   Abort any new incoming transaction  $t$  in the ordering phase
18: Upon  $b_{\text{last}}$  reaching  $W_H^n$ 
19:   Pause block formation thread until action  $a_{n+1}$  is derived
20:    $W_L^{n+1} \leftarrow W_H^n + \lfloor \Delta N_{\text{learn}} / |b_{n+1}| \rfloor$ 
21:    $W_H^{n+1} \leftarrow W_H^n + \lfloor \Delta N_{\text{episode}} / |b_{n+1}| \rfloor$ 
22:    $T^{n+1} \leftarrow T^n + \lfloor \Delta N_{\text{episode}} / |b_{n+1}| \rfloor \times |b_{n+1}|$ 
23:    $n \leftarrow n + 1$ 
24:   Carry out action  $a_{n+1}$ 
25:   Reset timer  $\tau$ 

```

for the next episode $n + 1$ from its local blockchain ledger (lines 1-3). Although most dimensions of the state are naturally consistent across different servers, there can be slight variations on the execution delay, e_i^{n+1} , and measured throughput. Thus, each server i multicasts a checkpoint message consisting of e_i^{n+1} and p_i^n to all other servers (line 4). AdaChain relies on the leader server to (1) collect a quorum Q of $2f + 1$ checkpoint messages, (2) calculate the median of observed throughput values to be the global reward p^n , and (3) calculate the median of the execution delay values e^{n+1} to be part of the global state (lines 5-7). Taking the median value prevents malicious servers from disrupting the predictive model’s training and inference by multicasting adversarial e_j^{n+1} and p_j^n that are abnormally high (or low). Once both values are computed, the leader multicasts a propose message, including the values and the set C of $2f + 1$ received checkpoint messages to all servers (line 8). Sending set C inside the message enables servers to validate e^{n+1} and p^n values. This is necessary because a malicious leader might compute the values incorrectly. Upon receiving the propose message, each server validates the message and multicasts an accept message to all other servers (lines 9-11). The goal of the accept phase is to prevent a malicious leader from sending different values of e^{n+1} and p^n to different servers. Each server then waits for $2f + 1$ matching accept messages before sending a commit message (lines 12-13). The accept and commit phases, similar to prepare and commit phases of PBFT, ensure that values are correct and replicated on a sufficient number of nodes. Finally, when a server receives $2f + 1$ commit messages, the predictive model will derive action a_{n+1} as described in Section 4 (lines 14-15). Note that since accept and commit messages are broadcast to all servers, even if a server has not received the propose message from the leader (due to the asynchronous nature of the network or the maliciousness of the leader), the server still has access to the values.

Algorithm 2 Slow path

▷ On each server i

- 1: **if** timer τ timeouts and b_{last} has not reached W_L^n **then**
- 2: pause block formation thread after committing the current block
- 3: Record performance p_i^n
- 4: Multicast $\langle \text{S-CHECKPOINT}, n, i, b_{\text{last},i}^n \rangle_{\sigma_i}$ to all servers

▷ On each server j where τ has not been expired

- 5: **Upon receiving** $f + 1$ valid S-CHECKPOINT messages from different servers
- 6: pause block formation thread after committing the current block
- 7: Record performance p_j^n
- 8: Multicast $\langle \text{S-CHECKPOINT}, n, j, b_{\text{last},j}^n \rangle_{\sigma_j}$ to all servers

▷ On the leader server l

- 9: **Upon receiving** valid S-CHECKPOINT messages from a quorum Q of $2f + 1$ servers
- 10: Compute $W_H^n \leftarrow \max\{b_{\text{last},j}^n \mid j \in Q\}$
- 11: Multicast $\langle \langle \text{S-PROPOSE}, W_H^n \rangle_{\sigma_l}, C' \rangle$ to all servers

▷ On each server i

- 12: **Upon receiving** a S-PROPOSE message from the leader
- 13: **if** W_H^n is valid (based on C') **then**
- 14: Multicast $\langle \text{S-ACCEPT}, n, i, W_H^n \rangle_{\sigma_i}$ to all servers
- 15: **Upon receiving** valid matching S-ACCEPT from $2f + 1$ different servers
- 16: Multicast $\langle \text{S-COMMIT}, n, i, W_H^n \rangle_{\sigma_i}$ to all servers
- 17: **Upon receiving** valid matching S-COMMIT from $2f + 1$ different servers
- 18: Resume block formation thread
- 19: Extract features $f_i^{n+1} = (w_i^{n+1}, c_i^{n+1}, r_i^{n+1}, e_i^{n+1})$ from block W_H^{n-1} to b_{last}
- 20: Multicast $\langle \text{CHECKPOINT}, n, i, f_i^{n+1}, p_i^n \rangle_{\sigma_i}$ to all servers

▷ On the leader server l

- 21: **for** every transaction t in the ordering phase **do**
- 22: $t.\text{episode} \leftarrow n$

▷ On each server i

- 23: **for** every transaction t committed by consensus **do**
- 24: **if** $t.\text{episode} \neq n$ **then**
- 25: abort t

In order to be responsive to workload changes, each episode is marked by the completion of $\lfloor \Delta N_{\text{episode}} / |b_n| \rfloor$ blocks where $\Delta N_{\text{episode}}$, as discussed in Section 3, is a constant hyperparameter of the system (10,000 transactions in the current deployment), and $|b_n|$ denotes the block size in episode n . As a result, each episode processes $\lfloor \Delta N_{\text{episode}} / |b_n| \rfloor \times |b_n|$ transactions, including transactions invalidated in MVCC validation due to conflicts. Specifically, when the number of committed transactions *in consensus* reached T^n , AdaChain early aborts transactions in the ordering phase (i.e., no more transactions will be committed by the consensus protocol) until AdaChain transitions into the next episode (lines 16-17). In AdaChain, the block formation thread waits for transactions to be committed, cuts the block, possibly performs dependency graph construction, reordering, or execution according to the current architecture, and lastly, commits the block. Once the number of committed blocks reaches the high watermark, the block formation thread will be paused until action a_{n+1} is derived (lines 18-19). This ensures exactly $\lfloor \Delta N_{\text{episode}} / |b_n| \rfloor$ blocks are committed in episode n on different servers. Note that the learning phase (including feature extraction, exchanging measurements, training, and inference) is triggered by low watermark W_L^n , and AdaChain keeps processing transactions using architecture a_n between W_L^n and W_H^n . Thus, as shown in Section 6.5, in most cases, architecture a_{n+1} is derived before reaching W_H^n , ensuring the high throughput of AdaChain.

5.2 Slow Path

Before AdaChain converges to the optimal architecture, the learning agent might occasionally choose “bad” architectures. The bad architectures might result in a high fraction of transactions being invalidated, or a slow growth of committed blocks (e.g., choosing

OX when the workload is highly compute-intensive, or choosing XOV+reorder when the contention is extremely high). In terms of wall-clock time, AdaChain should not be stuck in either scenario. While the normal path is capable of handling the first scenario, we further introduce a slow path to handle the scenario where the growth of committed blocks is slow.

Algorithm 2 presents the slow path operations. Note these are additional operations upon Algorithm 1. When the timer of server i timeouts and the index of the last committed block, b_{last} , has not reached the low watermark, server i pauses the block formation thread after committing the current block, records the performance p_i^n in the current episode, and multicasts a s-checkpoint message including the b_{last} to all other servers (lines 1-4). If a server j receives s-checkpoint messages from at least $f + 1$ servers, even if its timer has not expired, it pauses the block formation, records its performance, and multicasts a s-checkpoint message to all other nodes (lines 5-8). Since at most f Byzantine servers might send s-checkpoint messages maliciously, $f + 1$ messages are needed.

When the leader receives s-checkpoint messages from a quorum of $2f + 1$ servers, it finds the maximum index of the last committed block across all servers, W_H^n , and multicasts a s-propose message including W_H^n and the received $2f + 1$ s-checkpoint messages to all servers (lines 9-11). All servers validate the received s-propose message before two rounds of s-accept and s-commit communication, as shown in lines 12-16 (similar to the normal path). Each server then uses W_H^n as its high watermark and then resumes the block formation thread (lines 17-18). This ensures that in a slow path, the same number of blocks are committed in episode n across different servers. These operations might be expensive on the normal path, but are negligible on the slow path, compared to the timeout (15s in our case) and the poor performance before timeouts. The worst case happens when a fast server has not sent a s-checkpoint message, or its message has not been considered in the leader’s calculation of W_H^n . In this case, if the index of its last committed block is higher than W_H^n , the server needs to rollback those exceeding blocks. Similar to the normal path, each server also needs to exchange state and performance measurements to derive action a_{n+1} for the next episode (lines 19-20). Upon receiving transaction t for ordering at the leader l , the leader tags t with the current episode n as part of the sequence number (lines 21-22). When a server receives transactions committed by consensus protocol, it aborts transactions whose tagged episode is not equal to the current episode (lines 23-25). This ensures episode independence, i.e., there are no pending blocks across episodes in AdaChain. As a result, a bad architecture that triggers the slow path will not affect the performance of future episodes with promising architectures.

The normal path and slow path of AdaChain ensure two properties. First, transactions are strong serializable, and second, the world state is eventually consistent across different servers.

5.3 Security Analysis

AdaChain introduces two potential vulnerabilities that existing blockchain architectures do not have. First, malicious peers might attack the switching protocol. As illustrated in Sections 5.1 and 5.2, the protocol is robust to equivocation and poisoning attacks. Second, a malicious peer can manipulate its local model to choose a

different architecture from what honest peers choose within the same episode. Below, we briefly show that in this scenario, malicious peers cannot affect the correctness of honest peers.

Assume $f = 1$ and there are $3f + 1$ peers in the system, where P_1, P_2, P_3 are honest and P_4 is malicious. Without loss of generality, we assume the honest peers choose the XOV architecture. We first discuss the cases when P_4 is currently the leader in the consensus protocol. If P_4 chooses OX in the same episode, it will send out transaction proposals instead of endorsements, so honest peers can detect this mismatch and elect a new leader. If P_4 chooses XOV but with a batch size b_4 that is different from honest peers, it does not affect the correctness of honest peers. If P_4 chooses XOV with different batch sizes and uses different batch sizes for different peers (e.g., $b_{4,1}$ for P_1 , $b_{4,2}$ for P_2 , etc.), the honest peer will detect and reject these batches during consensus on batches and select a new leader. If P_4 chooses XOV but with an opposite reordering choice, since reordering happens locally on each peer according to its local model, the malicious peer cannot corrupt honest peers. In the cases where P_4 is not the leader in consensus protocol, the honest leader can detect type mismatch for messages originating from P_4 and discard them, while other honest peers work as normal.

6 EVALUATION

Our evaluation aims to answer the following questions:

- (1) Can AdaChain converge to the optimal architecture under a static workload without prior knowledge? (Section 6.2)
- (2) How well does AdaChain perform compared to existing fixed blockchain architectures when the workload changes? (Section 6.3)
- (3) How does the hardware setup (e.g., the type of CPU, network latency and bandwidth, etc.) affect the performance of AdaChain and existing blockchain architectures? (Section 6.4)
- (4) What overhead does AdaChain introduce? (Section 6.5)

6.1 Experimental Setup

We have implemented a prototype of AdaChain in C++ and Python. The blockchain peers which process transactions and carry out architecture switching are implemented in C++. We use gRPC for communications between peers and LevelDB [1] for storing the world states. The learning agents are implemented separately in Python due to its mature machine learning libraries. Each peer communicates with its local learning agent through gRPC.

Testbed. Our testbed consists of 4 c6220 bare-metal machines on CloudLab [29], each with two Xeon E5-2650v2 processors (8 cores each, 2.6Ghz), 64GB RAM(8 x 8GB DDR-3 RDIMMs, 1.86Ghz) and two 1TB SATA 3.5" 7.2K rpm hard drives. These machines are connected by two networks, each with one interface: (1) a 1 Gbps Ethernet control network; (2) a 10 Gbps Ethernet commodity fabric. Unless otherwise specified, we use the second network for all communication. We set the size of the execution thread pool equal to the number of cores on each peer.

System configuration. We run a single blockchain channel that consists of 3 peers on 3 different servers. As mentioned in Section 3, each peer in AdaChain serves as an executor as well as an orderer. The choice of consensus protocol in the ordering phase is configurable inside AdaChain, and we use Raft [48] for consistency

Workload	P_w	P_{hot}	N_{hot}	N_{trans}	T_{fire}	$T_{compute}$
A	0.2	0.95	5	300	50ms	5ms
B	0.5	0.99	10	100	50ms	1ms
C	0.5	0.1	10	300	50ms	10ms
D	0.9	0.95	1	100	50ms	0ms
E	0.5	0.99	10	100	50ms	5ms

Table 3: Specific workload parameters.

with Hyperledger Fabric and its variants. We run the client on a separate server with 3 threads, each firing transaction proposals to one specific peer. The reported throughput only considers *effective* transactions, i.e., excluding early aborted and invalidated transactions. Throughout this paper, we parameterize the architecture switching protocol as follows: normal path timeout is set to 15s, the low watermark is set to 7500 transactions, and the high watermark to set to 10000 transactions.

Workloads. To capture the diversity in real-world blockchain transactions, we implement a benchmark driver above SmallBank [28] to derive customized workloads with tunable parameters. The benchmark driver preloads the blockchain with 10k users, each with two accounts. We set N_{hot} of them as hot accounts. When firing transactions, the client randomly picks one of the five modifying transactions with probability P_w and the read-only transactions with probability $1 - P_w$. Each transaction has a certain probability to access the hot accounts, as controlled by the P_{hot} parameter. The client continuously fires N_{trans} transactions every T_{fire} milliseconds. To simulate computation-heavy transactions, each transaction has a $T_{compute}$ interval after it fetches the required world states from the key-value store and before it carries out subsequent operations.

We use workloads A-E throughout this paper, where workloads A-D are the same as in Figure 1. The specific parameters of workloads A-E are listed in Table 3. Unlike workloads A-D, the additional workload E is introduced later in the section that explores adaptivity to different hardware settings. Although workload E has only slight deviation from workload B, it renders blockchain architectures extremely sensitive to hardware setup (details in Section 6.4). Note that we have written our own benchmark driver because no existing benchmark captures all these variations in workloads.

6.2 Convergence under Static Workloads

Our first set of experiments aims to demonstrate that *AdaChain can rapidly converge to the optimal architecture under a static workload with no prior experiences required*. We run AdaChain for 100 episodes on four representative workloads (i.e., A-D). To compare AdaChain against the *a priori* optimal architecture, we also perform a grid search in the action space to find the optimal architecture for each workload. For our four workloads, we compare AdaChain with the four optimal static architectures.

Figure 3 plots the performance of AdaChain and four baselines for each workload, where the baseline curves are smoothed for better readability. The curve for AdaChain is not smoothed. Each workload has a different optimal architecture. For instance, XOV+reorder is optimal for workload A, suboptimal for workloads B and C, but the worst for workload D. Unlike a fixed architecture that cannot adjust itself even under a static workload, AdaChain always converges to the optimal architecture for each workload quickly within




Figure 3: Convergence of AdaChain to the optimal architecture under static workloads. Each plot shows the performance of AdaChain, XOV+reorder, XOV, OXII, and OX. While different blockchain architectures are optimal for each workload, AdaChain always reaches near-optimal performance.

Workload	XOV + reorder	XOV	OXII	OX	AdaChain
A	1532	1415	968	194	1425
B	897	866	1545	861	1426
C	3228	3235	940	98	3153
D	1	272	1494	1498	1447
Average	1414	1447	1237	663	1862
Worst	1	272	940	98	1425

Table 4: Effective throughput (tps) for each architecture in the last 20 episodes of each workload. AdaChain provides the best average and worst-case throughput across the four workloads.

40 episodes, no matter how bad the first episode (initial architecture) is. Due to Thompson sampling, AdaChain still performs some exploration in the architecture space even after convergence, as identified by the drops in the performance plot. Although these explorations do not seem useful under static workloads, they are crucial for finding optimal architectures under a changing workload which is more realistic in today’s BaaS environment (see Section 6.3 for performance improvements with exploration).

Comparing to exhaustive grid search (not shown in Figure 3) that takes n_a (the size of action space, i.e., 100 in our case) to converge and performs pure random exploration at all times, AdaChain converges much faster and strikes a better balance between exploiting known good actions and exploring unknown actions.

Average performance. AdaChain obviously does not outperform the optimal action in any workload (it is the optimal action, after all). However, we show that AdaChain does offer good average and worst-case performance after convergence. Table 4 shows the throughput for each blockchain on each workload in the last 20 episodes of execution. Even with a few performance drops due to exploration, AdaChain achieves both the best average throughput and the best worst-case throughput across all four workloads.

6.3 Adaptivity under a Changing Workload

Our next experiment focuses on the key benefits of AdaChain: *when the workload is changing, AdaChain can commit significantly more transactions than the best baseline during the same deployment period.* To emulate a changing workload, we run workload A for the first 15 minutes, followed by workloads B, C, D, and A, each for another 15 minutes. We use the same four baselines as in Figure 3, which are the optimal architectures for workload A-D when they are static.

Figure 4 shows the number of cumulative committed transactions with respect to time. During the entire 75 minutes, AdaChain




Figure 4: Cumulative committed transactions with time for a changing workload, showing AdaChain’s ability to maintain superior performance during workload shifts.

successfully completed 7.73×10^6 committed transactions, while the best baseline XOV completed 6.60×10^6 committed transactions. The worst baseline OX only completed 0.87×10^6 committed transactions. AdaChain can successfully commit 1.12 million (17%) more transactions than the best baseline during 75 minutes. The trend in Figure 4 also suggests the improvement of AdaChain would become increasingly significant with a longer deployment time and more variations in the workloads, which are common in today’s BaaS.

Interestingly, Figure 4 also shows a “catastrophic” effect for certain fixed architectures when the workload is changing. For instance, when transitioning back to workload A again (60-75 min), the slope of XOV+reorder is near zero, indicating poor performance where few if any transactions are completed successfully. However, if we start running XOV+reorder right from the beginning under workload A without any changes to the workload (Figure 3(a)), XOV+reorder would be the optimal architecture. The reason XOV+reorder performs poorly in workload D (45-60min) is due to the high overhead of Johnson’s algorithm with a large number of cycles, which slows down the block formation. Since the block formation is sequential, the number of pending blocks grows significantly. Thus, incoming transactions simulate on stale data and would fail in the MVCC validation phase, even when transitioning back to workload A again. OX suffers from similar problems due to a large number of pending blocks. This phenomenon also justifies our watermark-based design of AdaChain, as elaborated in Section 5.




Figure 5: Effective throughput of AdaChain in each episode. Here, red vertical line indicates when the workload shifts. The workload shifts every 15 minutes. The number of episodes per 15 minutes duration varies depending on the transactions arrival rate and compute intensity of workload.

To further investigate how AdaChain switches its architecture under a changing workload, we plot AdaChain’s effective throughput in each episode during the 75 minutes in Figure 5. The red dashed vertical line indicates when the workload shifts. Although workloads A and B have the same duration in terms of wall clock time (15 min), they vary in terms of the number of episodes. This is because, depending on the transaction arrival rate and compute intensity of different workloads, each episode (which is marked by the high watermark) may have a different time duration.

Figure 5 shows that when workloads shifts, AdaChain is able to quickly converge and perform competitively with the optimal architecture. For instance, when transitioning from workload A to B, AdaChain quickly converges to the new optimal (i.e., OXII) and achieves a 1450 tps throughput. In contrast, while XOV+reorder is optimal under workload A, as shown in Figure 3(b), it is able to reach only 900 tps when processing workload B, even in the best-case scenario where the catastrophic effect is avoided by starting with workload B and XOV+reorder right at the beginning. When transitioning from workload B to C, AdaChain quickly converges to the new optimal (i.e., XOV) and achieves a 3250 tps throughput. In comparison, OXII, which is optimal under workload B, is able to achieve only 980 tps under workload C (Figure 3(c)).

Due to Thompson sampling, AdaChain maintains some degree of exploration in the architecture space even after convergence, so as to avoid getting stuck at the local optimum. As AdaChain gains more experiences (i.e., data points) on a certain workload, the extent of exploration decreases, which is indicated by the less frequent drops within each 15 minutes period. Also, when AdaChain encounters a workload it has seen before (e.g., transition to workload A again in the last 15 minutes), AdaChain converges much faster than the first time and has less variation in performance.

6.4 Adaptivity under Different Hardware

Our next set of experiments demonstrates another operational benefit of AdaChain: *when deployed on different hardware configurations, AdaChain can rapidly converge to the optimal architecture for that hardware without manually re-configuring the blockchain architecture.* We use workload E on three different hardware setups: HW1 stands for single data center deployment, where the network connecting peers has low latency (0.15 ms) and high bandwidth (10 Gbps), and each peer has 16 CPU cores; HW2 also stands for single

	featurization	communication	training	inference	episode
mean	0.11s ± 0.01	0.14s ± 0.04	0.21s ± 0.04	0.01s ± 0.01	3.67s ± 2.12
median	0.11s	0.14s	0.21s	0.01s	2.57s
max	0.20s	0.27s	0.32s	0.02s	16.35s

Table 5: Overhead of each stage in AdaChain.

data center deployment, but with 2 CPU cores per peer; HW3 stands for a multi-data center deployment, with high latency (50 ms) and low bandwidth (1 Gbps) network, and 2 CPU cores per peer. As mentioned in Section 6.1, for HW1 and HW2, we use the commodity network fabric; for HW3, we use the control network as well as Linux netem [30] to inject delay to the NIC. For each hardware setup, we also perform a grid search to find the optimal architecture for that hardware: for HW1, the optimal is (OXII, blocksize = 100); for HW2, the optimal is (XOV, blocksize = 1), for HW3 the optimal is (OXII, blocksize = 100) again. Figure 6 plots AdaChain’s performance in each episode on HW1-HW3, along with the averaged performance of the optimal architecture for comparison.

As shown in Figure 6, even under the same workload, the hardware setup affects the effective throughput and thus affects the choice of the best architecture. For instance, OXII performs well when each server has enough compute resources (HW1 best arch in Figure 6(a)), but suffers when the servers have low compute resources (HW1 best arch in Figure 6(b)); StreamChain can perform well in a single data center deployment (HW2 best arch in Figure 6(b)), but suffers from the high round trip time when deployed across multiple data centers due to the small batch size it uses in the consensus protocol (HW2 best arch in Figure 6(c)). No matter what type of hardware AdaChain is deployed on, AdaChain can adapt itself to the optimal architecture for that hardware.

More importantly, for any kind of unseen hardware setup, users of AdaChain do not need to recollect data and retrain the machine learning model offline. AdaChain is an online system that learns from its past experiences and balances exploitation and exploration. With AdaChain, BaaS can have humans completely out of the loop.

6.5 Overhead of Learning

Our last experiment evaluates the additional overhead incurred by AdaChain’s learning framework. We repeat the experiment in Section 6.3 and profile every stage that involves the learning agent. We report the results along with the episode duration in Table 5.




Figure 6: Convergence of AdaChain to the optimal architecture under different hardware.

Before deriving the architecture for the next episode, AdaChain needs to go through feature extraction, communication, training and inference phases in sequence (Section 3). Table 5 shows that all of these four phases have low average overhead with low variance, as compared to the episode duration. The additional mean overhead time (0.47s) is only 12.8% of the average episode duration (3.67s).

More importantly, the 12.8% overhead can be masked by parallelizing transaction processing and learning. The learning phase only occurs between the low and high watermark period, which constitutes 25% duration with each episode. Considering that the median episode duration is 2.57s, this interval time is 0.64s which is higher than the mean overhead of 0.47s. During this interval, the peers continue to process transactions based on the current architecture, while in parallel, the learning agent goes through the four stages to determine the architecture for the next episode. Such parallel execution ensures that the overhead of AdaChain does not adversely affect its effective throughput, as long as there are some spare CPU cycles devoted to the learning agent.

Finally, we note that the episode duration has a high standard deviation due to two reasons: (1) AdaChain triggers the slow path when making poor architecture choices; (2) even in the normal path, the wall-clock time needed to reach a high watermark depends on the workload, e.g., the episode duration is longer under a compute-intensive or low-load workload.

Unlike deep neural networks, which are especially expensive to train, the random forest model used by AdaChain has moderate training overhead. With thousands of data points in the experience buffer, AdaChain only incurs a maximum training overhead of 0.32 seconds. Moreover, the succinct action space design also results in a lightweight inference phase, i.e., maximum 0.02 seconds. When AdaChain is deployed for a long run that lasts for months or years, techniques such as periodic resampling and limiting the length of experience buffer [46] can be utilized.

7 RELATED WORK

In this section, we briefly survey several related research lines.

Permissioned blockchains. A permissioned blockchain system consists of a set of known, identified, but possibly untrusted participants. Permissioned blockchains have been analyzed in different surveys and empirical studies [12, 19, 22, 27, 28, 53, 57, 57]. Over the past few years, several benchmarks have been proposed to facilitate studying the performance of blockchains. Hyperledger Caliper [6]

benchmark framework is proposed to evaluate blockchain systems developed within the Hyperledger project. Blockbench [28], on the other hand, is able to benchmark all permissioned blockchains. Using Blockbench, blockchain systems can be evaluated under an existing, e.g., YCSB or SmallBank, or a newly implemented benchmark. Chainhammer [42] is another benchmark tool that can be used to evaluate Ethereum-based blockchains under extremely high loads. Finally, Diablo-v2 [34] presents a unified framework consisting of 5 realistic Decentralized Applications and their corresponding workloads. AdaChain models similar workload characteristics as Diablo and is also able to support the combination of such workloads, e.g., a contentious compute-intensive workload.

Contextual multi-armed bandits. Contextual multi-armed bandits [64] and Thompson sampling [58] have both been studied extensively [32, 39, 49, 55]. Thompson sampling has also been applied in other database context, such as parameterized query optimization [60], query optimization [46], and cloud workload management [47].

Learned systems. More generally, many recent works have applied machine learning concepts to systems components. These works, falling under the umbrella of machine programming [33], cannot be exhaustively enumerated here, but we refer the reader to past work on indexing [41], cardinality estimation [40, 44], index selection [26], database tuning [50], scheduling [45], garbage collection [21], and concurrency control for in-memory databases [61].

8 CONCLUSION

In this paper, we presented AdaChain, an adaptive blockchain framework that leverages reinforcement learning to dynamically switch between different blockchain architectures based on the workload. AdaChain is able to identify the optimal blockchain architecture as workload changes, obtaining significantly higher throughput compared to fixed architecture settings. AdaChain requires low additional overhead, which can be masked by parallelizing the transaction processing and learning phases.

As future work, we are exploring expanding the learning framework to cover other aspects of blockchain architectures, e.g., choosing the best performing consensus protocol. Another intriguing future direction is to figure out whether our learning framework can be used to uncover new effective architectures not previously explored by human experts.

REFERENCES

- [1] [n. d.]. <https://github.com/google/leveldb>
- [2] [n. d.]. AWS. Amazon quantum ledger database (QLDB). <https://aws.amazon.com/qldb/>.
- [3] [n. d.]. Blockchain on AWS Enterprise blockchain made real. <https://aws.amazon.com/blockchain/>.
- [4] [n. d.]. Chain. <http://chain.com>.
- [5] [n. d.]. Corda. <https://github.com/corda/corda>.
- [6] [n. d.]. Hyperledger Caliper. <https://www.hyperledger.org/use/caliper>.
- [7] [n. d.]. Hyperledger Iroha. <https://github.com/hyperledger/iroha>.
- [8] [n. d.]. IBM Blockchain Platform. <https://www.ibm.com/cloud/blockchain-platform>.
- [9] [n. d.]. Oracle Blockchain. <https://www.oracle.com/blockchain/>.
- [10] Shipra Agrawal and Navin Goyal. 2013. Further Optimal Regret Bounds for Thompson Sampling. In *The International Conference on Artificial Intelligence and Statistics (AISTATS '13)*.
- [11] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2019. Par-Blockchain: Leveraging Transaction Parallelism in Permissioned Blockchain Systems. In *Int. Conf. on Distributed Computing Systems (ICDCS)*. IEEE, 1337–1347.
- [12] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2021. Permissioned Blockchains: Properties, Techniques and Applications. In *SIGMOD Int. Conf. on Management of Data*. 2813–2820.
- [13] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2021. SharPer: Sharding Permissioned Blockchains Over Network Clusters. In *SIGMOD Int. Conf. on Management of Data*. ACM, 76–88.
- [14] Mohammad Javad Amiri, Joris Duguépéroux, Tristan Allard, Divyakant Agrawal, and Amr El Abbadi. 2021. SEPAR: Towards Regulating Future of Work Multi-Platform Crowdfunding Environments with Privacy Guarantees. In *Proceedings of The Web Conf. (WWW)*. 1891–1903.
- [15] Mohammad Javad Amiri, Boon Thau Loo, Divyakant Agrawal, and Amr El Abbadi. 2022. Qanaat: A Scalable Multi-Enterprise Permissioned Blockchain System with Confidentiality Guarantees. *Proc. of the VLDB Endowment* 15, 11 (2022), 2839–2852.
- [16] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, et al. 2018. Hyperledger Fabric: a distributed operating system for permissioned blockchains. In *European Conf. on Computer Systems (EuroSys)*. ACM, 30:1–30:15.
- [17] Leo Breiman. 1996. Bagging Predictors. In *Machine Learning (Maching Learning '96)*.
- [18] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (Oct. 2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [19] Christian Cachin and Marko Vukolić. 2017. Blockchain Consensus Protocols in the Wild. In *Int. Symposium on Distributed Computing (DISC)*. 1–16.
- [20] Miguel Castro, Barbara Liskov, et al. 1999. Practical Byzantine fault tolerance. In *Symposium on Operating systems design and implementation (OSDI)*, Vol. 99. USENIX Association, 173–186.
- [21] Lujing Cen, Ryan Marcus, Hongzi Mao, Justin Gottschlich, Mohammad Alizadeh, and Tim Kraska. 2020. Learned Garbage Collection. In *Proceedings of the 4th ACM SIGPLAN International Workshop on Machine Learning and Programming Languages (MAPL @ PLDI '20)*. ACM. <https://doi.org/10.1145/3394450.3397469>
- [22] Jeeta Ann Chacko, Ruben Mayer, and Hans-Arno Jacobsen. 2021. Why Do My Blockchain Transactions Fail? A Study of Hyperledger Fabric. In *SIGMOD Int. Conf. on Management of Data*. ACM, 221–234.
- [23] Olivier Chapelle and Lihong Li. 2011. An empirical evaluation of Thompson sampling. In *Advances in neural information processing systems (NIPS'11)*.
- [24] JP Morgan Chase. 2016. Quorum white paper.
- [25] Sam Daley. 2021. 18 Blockchain-as-a-Service Companies Making the DLT More Accessible. <https://builtin.com/blockchain/blockchain-as-a-service-companies>.
- [26] Bailu Ding, Sudipto Das, Ryan Marcus, Wentao Wu, Surajit Chaudhuri, and Vivek R. Narasayya. 2019. AI Meets AI: Leveraging Query Executions to Improve Index Recommendations. In *38th ACM Special Interest Group in Data Management (SIGMOD '19)*.
- [27] Tien Tuan Anh Dinh, Rui Liu, Meihui Zhang, Gang Chen, Beng Chin Ooi, and Ji Wang. 2018. Untangling blockchain: A data processing view of blockchain systems. *IEEE transactions on knowledge and data engineering* 30, 7 (2018), 1366–1385.
- [28] Tien Tuan Anh Dinh, Ji Wang, Gang Chen, Rui Liu, Beng Chin Ooi, and Kian-Lee Tan. 2017. Blockbench: A framework for analyzing private blockchains. In *SIGMOD Int. Conf. on Management of Data*. ACM, 1085–1100.
- [29] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, et al. 2019. The Design and Operation of {CloudLab}. In *2019 USENIX annual technical conference (USENIX ATC 19)*. 1–14.
- [30] The Linux Foundation. 2021. netem. Retrieved July 4, 2022 from <https://wiki.linuxfoundation.org/networking/netem>
- [31] Zerui Ge, Dumitrel Loghini, Beng Chin Ooi, Pingcheng Ruan, and Tianwen Wang. 2022. Hybrid blockchain database systems: design and performance. *Proceedings of the VLDB Endowment* 15, 5 (2022), 1092–1104.
- [32] Aditya Gopalan, Shie Mannor, and Yishay Mansour. 2014. Thompson Sampling for Complex Online Problems. In *International Conference on Machine Learning (ICML '14, Vol. 14)*. 100–108.
- [33] Justin Gottschlich, Armando Solar-Lezama, Nesime Tatbul, Michael Carbin, Martin Rinard, Regina Barzilay, Saman Amarasinghe, Joshua B. Tenenbaum, and Tim Mattson. 2018. The three pillars of machine programming. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages (MAPL 2018)*. Association for Computing Machinery, Philadelphia, PA, USA, 69–80. <https://doi.org/10.1145/3211346.3211355>
- [34] Vincent Gramoli, Rachid Guerraoui, Andrei Lebedev, Chris Natoli, and Gauthier Voron. 2022. *Diablo-v2: A Benchmark for Blockchain Systems*. Technical Report.
- [35] Gideon Greenspan. 2015. MultiChain private blockchain-White paper. URL: <http://www.multichain.com/download/MultiChain-White-Paper.pdf> (2015).
- [36] Hyperledger. [n. d.]. Private Data Collections: A High-Level Overview. <https://www.hyperledger.org/blog/2018/10/23/private-data-collections-a-high-level-overview>.
- [37] Zsolt István, Alessandro Sorniotti, and Marko Vukolić. 2018. StreamChain: Do Blockchains Need Blocks?. In *Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers (SERIAL)*. ACM, 1–6.
- [38] Manos Kapritsos, Yang Wang, Vivien Quema, Allen Clement, Lorenzo Alvisi, Mike Dahlin, et al. 2012. All about Eve: Execute-Verify Replication for Multi-Core Servers.. In *Symposium on Operating systems design and implementation (OSDI)*, Vol. 12. USENIX Association, 237–250.
- [39] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. 2012. Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory (ALT '12)*.
- [40] Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter Boncz, and Alfons Kemper. 2019. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In *9th Biennial Conference on Innovative Data Systems Research (CIDR '19)*. <http://arxiv.org/abs/1809.00677>
- [41] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The Case for Learned Index Structures. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3183713.3196909>
- [42] Andreas Krueger. 2019. Chainhammer: Ethereum benchmarking. <https://github.com/drandreaskrueger/chainhammer>.
- [43] Jae Kwon. 2014. Tendermint: Consensus without mining. (2014).
- [44] Henry Liu, Mingbin Xu, Ziting Yu, Vincent Corvini, and Calisto Zuzarte. 2015. Cardinality Estimation Using Neural Networks. In *Proceedings of the 25th Annual International Conference on Computer Science and Software Engineering (CASCON '15)*. IBM Corp., Riverton, NJ, USA, 53–59. <http://dl.acm.org/citation.cfm?id=2886444.2886453>
- [45] Hongzi Mao, Malte Schwarzkopf, Shaileshh Bojja Venkatakrisnan, Zili Meng, and Mohammad Alizadeh. 2018. Learning Scheduling Algorithms for Data Processing Clusters. *arXiv:1810.01963 [cs, stat]* (2018). <http://arxiv.org/abs/1810.01963>
- [46] Ryan Marcus, Parimarjan Negi, Hongzi Mao, Nesime Tatbul, Mohammad Alizadeh, and Tim Kraska. 2021. Bao: Making Learned Query Optimization Practical. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21)*. China. <https://doi.org/10.1145/3448016.3452838> Award: 'best paper award'.
- [47] Ryan Marcus and Olga Papaemmanouil. 2017. Releasing Cloud Databases from the Chains of Performance Prediction Models. In *8th Biennial Conference on Innovative Data Systems Research (CIDR '17)*. San Jose, CA. tex.authors= Ryan Marcus and Olga Papaemmanouil.
- [48] Diego Ongaro and John K Ousterhout. 2014. In search of an understandable consensus algorithm. In *Annual Technical Conf. (ATC)*. USENIX Association, 305–319.
- [49] Ian Osband and Benjamin Van Roy. 2015. Bootstrapped Thompson Sampling and Deep Exploration. *arXiv:1507.00300 [cs, stat]* (July 2015). <http://arxiv.org/abs/1507.00300>
- [50] Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon, Todd C. Mowry, Matthew Perron, Ian Quah, Siddharth Santurkar, Anthony Tomic, Skye Toor, Dana Van Aken, Ziqi Wang, Yingjun Wu, Ran Xian, and Tieying Zhang. 2017. Self-Driving Database Management Systems. In *8th Biennial Conference on Innovative Data Systems Research (CIDR '17)*.
- [51] Zhe Peng, Cheng Xu, Haixin Wang, Jinbin Huang, Jianliang Xu, and Xiaowen Chu. 2021. P2B-Trace: Privacy-Preserving Blockchain-based Contact Tracing to Combat Pandemics. In *SIGMOD Int. Conf. on Management of Data*. 2389–2393.
- [52] Zhe Peng, Jianliang Xu, Xiaowen Chu, Shang Gao, Yuan Yao, Rong Gu, and Yuzhe Tang. 2021. Vfchain: Enabling verifiable and auditable federated learning via blockchain systems. *IEEE Transactions on Network Science and Engineering* (2021).
- [53] Pingcheng Ruan, Tien Tuan Anh Dinh, Dumitrel Loghini, Meihui Zhang, Gang Chen, Qian Lin, and Beng Chin Ooi. 2021. Blockchains vs. Distributed Databases: Dichotomy and Fusion. In *SIGMOD Int. Conf. on Management of Data*. 1504–1517.

- [54] Pingcheng Ruan, Dumitrel Loghin, Quang-Trung Ta, Meihui Zhang, Gang Chen, and Beng Chin Ooi. 2020. A Transactional Perspective on Execute-order-validate Blockchains. In *SIGMOD Int. Conf. on Management of Data*. ACM, 543–557.
- [55] Daniel Russo and Benjamin Van Roy. 2014. An information-theoretic analysis of Thompson sampling. *Journal of Machine Learning Research* (2014).
- [56] Ankur Sharma, Felix Martin Schuhknecht, Divya Agrawal, and Jens Dittrich. 2019. Blurring the lines between blockchains and database systems: the case of hyperledger fabric. In *SIGMOD Int. Conf. on Management of Data*. ACM, 105–122.
- [57] Man-Kit Sit, Manuel Bravo, and Zsolt István. 2021. An experimental framework for improving the performance of BFT consensus for future permissioned blockchains. In *Proceedings of the 15th ACM Int. Conf. on Distributed and Event-based Systems*. 55–65.
- [58] William R. Thompson. 1933. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika* (1933).
- [59] Feng Tian. 2017. A supply chain traceability system for food safety based on HACCP, blockchain & Internet of things. In *Int. Conf. on service systems and service management (ICSSSM)*. IEEE, 1–6.
- [60] Kapil Vaidya, Anshuman Dutt, Vivek Narasayya, and Surajit Chaudhuri. 2022. Leveraging query logs and machine learning for parametric query optimization. *Proceedings of the VLDB Endowment* 15, 3 (Feb. 2022), 401–413. <https://doi.org/10.14778/3494124.3494126>
- [61] Jia-Chen Wang, Ding Ding, Huan Wang, Conrad Christensen, Zhaoguo Wang, Haibo Chen, and Jinyang Li. 2021. Polyjuice: High-Performance Transactions via Learned Concurrency Control. In *OSDI*. 198–216.
- [62] Chenyuan Wu, Mohammad Javad Amiri, Jared Asch, Heena Nagda, Qizhen Zhang, and Boon Thau Loo. 2022. FlexChain: An Elastic Disaggregated Blockchain. *Proc. of the VLDB Endowment* 16, 01 (2022), 23–36.
- [63] Xinying Yang, Yuan Zhang, Sheng Wang, Benquan Yu, Feifei Li, Yize Li, and Wenyuan Yan. 2020. LedgerDB: a centralized ledger database for universal audit and verification. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3138–3151.
- [64] Li Zhou. 2016. A Survey on Contextual Multi-armed Bandits. *arXiv:1508.03326 [cs]* (Feb. 2016). <http://arxiv.org/abs/1508.03326>