


Efficient Parallel Execution of Blockchain Transactions Leveraging Conflict Specifications

Parwat Singh Anjana* ✉ 


Supra Research, Udaipur, India

Matin Amini ✉ 


University of Southern California, Los Angeles, CA, USA

Rohit Kapoor ✉ 

Supra Research, New Delhi, India

Rahul Parmar ✉ 

Supra Research, Ahmedabad, India

Raghavendra Ramesh ✉ 

Supra Research, Brisbane, Australia

Srivatsan Ravi* ✉ 

Supra Research, Los Angeles, CA, USA

University of Southern California, Los Angeles, CA, USA

Joshua Tobkin ✉

Supra Research, Taipei, Taiwan

Abstract

Parallel execution of smart contract transactions in large multicore architectures is critical for higher efficiency and improved throughput. The main bottleneck for maximizing the throughput of a node through parallel execution is transaction conflict resolution: when two transactions interact with the same data, like an account balance, their order matters. Imagine one transaction sends tokens from account A to account B, and another tries to send tokens from account B to account C. If the second transaction happens before the first one, the token balance in account B might be wrong, causing the entire system to break. Conflicts like these must be managed carefully, or you end up with an inconsistent, unusable blockchain state.

Traditional software transactional memory (STM) has been identified as a possible abstraction for the concurrent execution of transactions within a block, with Block-STM pioneering its application for efficient blockchain transaction processing on multicore validator nodes. This paper presents a parallel execution methodology that leverages conflict specification information of the transactions for block transactional memory (BTM) algorithms. Our experimental analysis, conducted over synthetic transactional workloads and real-world blocks, demonstrates that BTMs leveraging conflict specifications outperform their plain counterparts on both EVM and MoveVM. Our proposed BTM implementations achieve up to $1.75\times$ speedup over sequential execution and outperform the state-of-the-art Parallel-EVM (PEVM) execution by up to $1.33\times$ across synthetic workloads.

2012 ACM Subject Classification Computing methodologies → Parallel computing methodologies

Keywords and phrases Blockchain, Smart Contract, Parallel Execution, Conflict Specifications

Digital Object Identifier 10.4230/LIPIcs.AFT.2025.29

Related Version *Full Version:* <https://arxiv.org/abs/2503.03203>

*Corresponding authors.



© Parwat Singh Anjana, Matin Amini, Rohit Kapoor, Rahul Parmar, Raghavendra Ramesh, Srivatsan Ravi, and Joshua Tobkin;

licensed under Creative Commons License CC-BY 4.0

7th Conference on Advances in Financial Technologies (AFT 2025).

Editors: Zeta Avarikioti and Nicolas Christin; Article No. 29; pp. 29:1–29:26



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Blockchains are in a race to reduce transaction latency, improving user experience and increasing throughput to meet anticipated demand. This optimization effort has examined every step of a blockchain transaction, including data dissemination, ordering, execution, and storage. In this paper, we investigate the state-of-the-art in optimizing the execution step of a transaction workflow.

Blockchains require all nodes to deterministically reach the same final state when executing a block of transactions. To leverage multicore architectures, modern blockchain designs seek to maximize parallel execution throughput while preserving deterministic consistency across all nodes. Traditional software transactional memory (STM) has been identified as a possible abstraction for the concurrent execution of transactions within a block. However, the execution order must follow a *preset order* of transactions in a block. The main bottleneck for maximizing the throughput of a node through block transactional memory (BTM) that respects the preset order is transaction *conflict* resolution: when two transactions, T_1 and T_2 in a block, interact with the same state, like an account balance, order in which the reads and writes on the state occurs matters for the correctness of the execution. Imagine that a transaction T_1 sends tokens from account A to account B , and another transaction T_2 tries to send tokens from account B to account C . If T_2 happens before the first one, the token balance in account B might be wrong, causing the entire system to break. Conflicts like these must be managed carefully, or you end up with an inconsistent, unusable blockchain state.

Ethereum maintain consistency by sequentially executing a block of transactions, underutilizing the multicore architectures of its nodes. In contrast, Solana [33] and Aptos [5] utilize parallel execution, though in different settings. Aptos employs the classical STM technique to execute transactions in parallel with BTM, without requiring any priori read-write access specifications with transactions. Aptos' Block-STM [17], represents the current state-of-the-art, applying classical STM [32] techniques to the execution of ordered blocks of transactions. STM techniques are typically speculative or optimistic, meaning that they attempt to execute as many transactions in parallel as possible, while detecting and resolving conflicts. If a transaction reads a value that later changes due to another transaction, it must be re-executed to maintain consistency. In contrast, Solana's parallel execution requires transactions to be tagged with the accounts read and written, though this is not a STM-based algorithm. We wondered whether the read-write access specifications could enhance a BTM algorithm's performance, and our findings indicated a positive correlation.

Firstly, we observe that some popular blockchains equip transactions a priori with *read-write sets*, i.e., the sets of accounts that the transaction *may* access. For instance, Solana leverages user-provided read-write sets as access specifications in its lock-profile-based iterative, parallel execution strategy [38]. Similarly, Sui [16] utilizes user-provided access specifications to execute transactions in a causally-ordered manner. Secondly, we further observe that, when such access specifications are not available a priori, it is possible to derive access specifications for public entry functions statically at the time of deployment of smart contracts using data-flow analysis on the smart contract code. This is a one-time computation, enabling efficient utilization of these specifications during transaction execution. These access specifications essentially serve as *conflict specifications* for any parallel execution technique on ordered blocks of transactions. As a complementary approach, we observe that it is possible to build conflict specifications efficiently at runtime by checking whether the accounts affected by a set of transactions (for instance, *payment transfers* in Aptos Move or *ETH* or *ERC-20* transfers in Ethereum [12]) are mutually disjoint.

This is where we see the challenge of improving BTM to further optimize execution by leveraging conflict specifications. Specifically, we seek to answer the following questions: Do conflict specifications improve transaction-execution throughput in existing BTM techniques? If so, what is the best BTM technique to optimally leverage these conflict specifications? By constructively answering these questions, we aim to advance the frontiers of parallel transaction execution, driving significant improvements in blockchain scalability and efficiency.

Contributions. Our key contributions are as follows:

1. Leveraging our insight that conflict specifications must be exploited for maximizing throughput in parallel execution, we detail two BTM algorithms, dBTM and oBTM, that can efficiently execute block transactions when *sound*, but possibly *incomplete* conflict specifications are available.
2. We present implementations of our algorithms for Ethereum VM (EVM) and Aptos MoveVM and conduct a detailed empirical analysis on real-world blocks, thoroughly analyzing the best and worst-case performance. Additionally, we design a workload generator for analyzing performance on large blocks and unconventional workloads.
3. We present a rigorous formalism and implementation for how conflict specifications are efficiently constructed for EVM and MoveVM. We then present a conflict analyzer integrated implementation (iBTM) that outperforms existing parallel execution approaches.
4. We evaluated the performance of iBTM against sequential and state-of-the-art Parallel-EVM (PEVM) [28] on both synthetic and historical workloads. In synthetic workload, the proposed iBTM achieve up to $1.75\times$ speedup over sequential and $1.33\times$ over PEVM and maintain consistent improvements, with an average speedup of $1.24\times$ over PEVM. It achieves a maximum speedup of $2.43\times$ over sequential execution on historical workloads.

The methodology is rigorously demonstrated for EVM and MoveVM, for which conflict specifications are algorithmically derived as part of our integrated implementation. However, we remark that the purpose of our modular presentation and detailed ablation studies is to demonstrate how easily our parallel execution algorithms can be applicable to other blockchain ecosystems, specifically Sui and Solana, which explicitly provide a priori transactional access sets that can be used to derive conflict specifications. By identifying dependencies upfront, we can avoid unnecessary rollbacks and retries, making our system much more efficient for prevalent blockchain infrastructures. This allows the system to automatically separate transactions that need special handling due to dependencies, making the entire execution process more streamlined and as we demonstrate, highly efficient. Finally, we present how this approach also unleashes the possibility of workload *adaptive* execution that can leverage a sequential or parallel execution algorithm, depending on the block’s “conflict threshold”.

Roadmap. The rest of the paper is organized as follows. Sections 2 and 3 presents related work and motivation on parallel execution of block transactions, respectively. Section 4 introduces the system model and the proposed BTM algorithms. Section 5 presents experimental results on parallel execution using conflict specifications. Section 6 formalize and implement conflict specification generation and evaluate the full integration in both real-world and synthetic workloads. For succinctness of presentation, we focus on the EVM in the paper and only summarize the MoveVM evaluation results. We conclude in Section 7.

2 Related Work

This section reviews smart contract execution models and existing parallel execution algorithms for block transactional execution.

Several models have been proposed for the execution of transactions in blockchains; one such model involves the *block proposer* executing transactions, generating a block containing state differences, and subsequently propagating this block across the network for validators to validate. Another approach entails the block proposer determining the transaction order in a block, after which all nodes reach consensus on this order before executing transactions in parallel. The former is known as the *Ethereum model* [12], while the latter is known as the *Aptos model* [5]. Another model facilitates parallel execution by incorporating read-write sets with transactions; this is commonly called the *Solana model* [33]. A fourth model exploits resource ownership to enable parallel execution, a methodology known as the *Sui model* [36].

These four execution models can be broadly categorized into two distinct classes that aim to optimize transaction execution. The first class is called the *read-write aware execution* which relies on transaction access hints provided by clients to facilitate parallel execution either through preprocessing in the form of a directed acyclic graph (DAG) or runtime techniques based on lock profiling. The second class comprises techniques that leverage run-time execution techniques, such as STM or locks, to optimistically execute transactions. We call this approach the *read-write oblivious execution*.

Read-write Aware Execution. Transactions on blockchains such as Solana [33] and Sui [36] upfront specify the accounts they access in read and write mode during execution. These access hints are used to enable parallel execution, either through static analysis or by employing a runtime scheduler that resolves read-write set conflicts and executes transactions in parallel. Alternatively, a DAG can be constructed from the read-write access sets to partition transactions into independent groups (iterations in Solana) for parallel execution. Specifically, Solana’s [33] SeaLevel [37, 38], leverages read-write sets and lock profiling to execute transactions iteratively. Each iteration involves a locking phase to detect conflicts and an execution phase where non-conflicting transactions are executed in parallel, while conflicting transactions are deferred to subsequent iterations until all transactions in the block are processed. The iteration information is then included in the block by the block proposer to support parallel execution during validation at the validators.

In contrast, Sui [36] introduces an object-based state model to identify independent transactions. Objects are shared or exclusively owned, each with a unique identity and owner address. Dependency identification is simplified by determining whether multiple transactions access the same shared object. Based on this model, Sui enables parallel execution [15, 16] where transactions on owned objects are executed in parallel and can completely bypass consensus, while those on shared objects run sequentially through consensus to avoid conflicts. Other works, such as ParBlockchain [1] and Hyperledger Sawtooth [30], employ lock-based or conflict-analysis techniques, while DiPETrans [6] and the efficient scheduler for Sawtooth [25] utilize DAG-based read-write aware execution.

Read-write Oblivious Execution. In this class, the goal is to execute an ordered set of transactions in parallel as if they were executed sequentially and arrive at the same state. The key idea is that some transactions may not be conflicting, i.e, they do not read or write any common states; hence, can be executed in parallel, enabling execution acceleration that arrives at the correct sequential result. For the Ethereum model, Dickerson et al. [11] propose, the first pioneering work on parallel execution of block transactions. They proposed using pessimistic ScalaSTM for parallel execution at the block proposer and a happen-before graph for parallel execution at the validators. Later, Anjana et al. [2, 3] proposed an optimistic STM (OSTM) based multi-version timestamp ordering protocol for parallel execution at

the proposer, while a DAG-based efficient parallel execution at validators. Saraph and Herlihy [29] proposed a simple *bin-based two-phase* approach. In the first phase, the proposer uses locks and tries to execute transactions concurrently by rolling back those that lead to conflict(s). Aborted transactions are kept in a sequential bin and executed sequentially in the second phase. Later, OptSmart [4] proposed an approach that combines the idea of bin-based approach with the OSTM approach for efficient parallel execution.

In the Aptos model, differed execution of transaction based on preset serializable and multi-version concurrency control is proposed in the Block-STM [17]. Rather than speculatively executing block transactions in any order, they employ it on ordered-sets, called the *preset order*. Each validator uses Block-STM independently to execute a leader proposal of an ordered set of transactions in parallel to get the same state. This has currently been implemented on the Aptos blockchain [5] and is the most promising approach, as it does not require additional information to be attached to the transactions or in the block for parallel execution. It has been adopted for execution on the Polygon PoS Chain [26], where it is already live on the mainnet. The proposer uses Block-STM to execute transactions in parallel and includes a DAG in the block to allow deterministic and safe parallel execution at validators. Recently, SeiGiga [31] has also adopted Block-STM.

3 Motivation and Overview

This section motivates the need for parallel execution in blockchain systems and outlines the benefit of leveraging conflict specification for parallel execution.

Motivation. Naturally, the throughput of STM-based execution varies widely based on the level of transaction conflicts, aborts, and re-executions (triggered re-validations of the transaction's read state), presenting both a challenge and an opportunity to further optimize.

Consider the scenario shown in Figure 1, in which we have two transactions T_1 and T_2 running concurrently. We consider that T_1 must precede T_2 in the preset order (denoted $T_1 \rightarrow T_2$). Without *a priori* knowledge of read-write conflicts in the set of accounts (states)

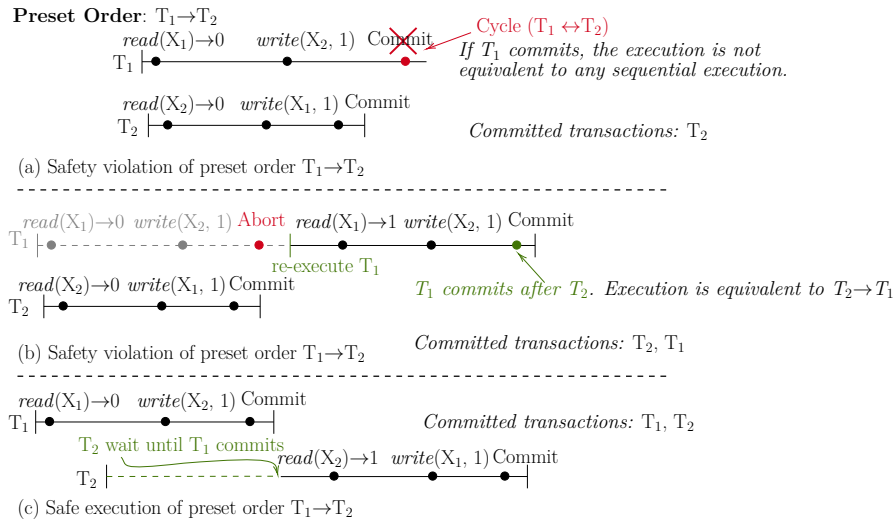
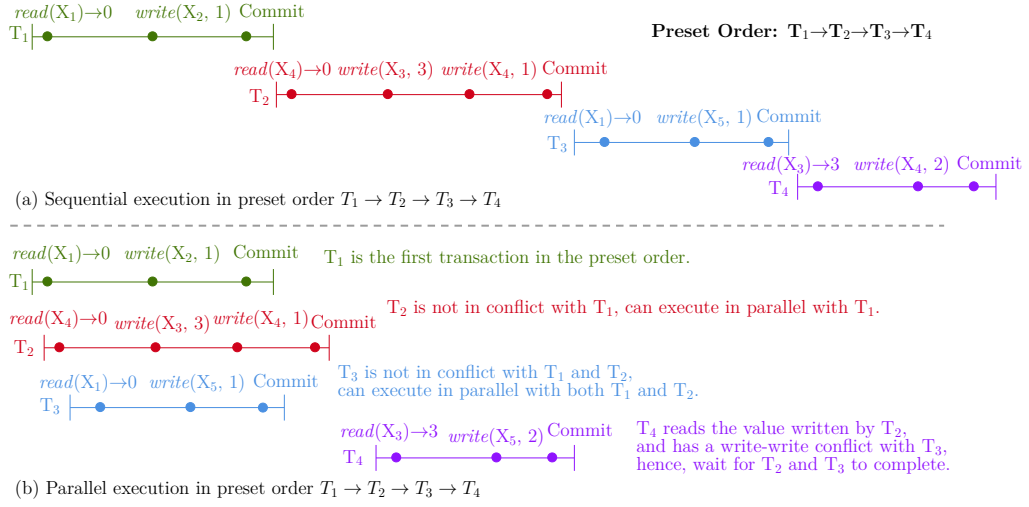


Figure 1 Safe execution in preset order: (a) and (b) show unsafe executions due to arbitrary or incorrect serialization; (c) shows a safe execution where T_2 waits for T_1 to commit, preserving the preset serialization order.



■ **Figure 2** Leveraging conflict specifications a priori for parallel execution.

accessed, committing T_2 prior to committing T_1 may result in a safety violation. To illustrate this, consider the following execution: T_1 reads an account X_1 (value 0), following which T_2 reads X_2 (value 0), followed by a write of a new value 1 to X_1 . Observe that if T_2 commits at this point in the execution and if T_1 writes a new value 1 to X_2 after the commit of T_2 , then the resulting execution does not respect the preset order in any extension. Clearly, if T_1 commits, the resulting execution is not equivalent to any sequential execution, as shown in Figure 1a. Alternatively, if T_1 aborts and then re-starts, any read of X_2 will return the value 1 that is written by T_2 , thus not respecting the preset order, as illustrated in Figure 1b. Consequently, the only possible way to avoid this, requires T_2 to wait until T_1 commits (see Figure 1c). Now consider a modification of this execution in which T_2 reads X_3 and writes to X_4 . In this case, T_2 does not have a *read-from conflict* with T_1 allowing T_1 and T_2 to run in parallel with almost no synchronization. However, this is possible only if threads that execute these transactions are *aware* of the conflict prior to the execution.

After studying both read-write oblivious and aware models (taxonomized in Section 2), we conclude that the approach that leverage conflict specifications is the best way forward. It can be considered the “Goldilocks” approach which creates conflict specifications to overcome the limitations of parallel execution imposed by the preset order of transaction. This is what will allow us to scale with parallel execution efficiently and maximize throughput.

Overview. Consider the execution of four transactions T_1 – T_4 depicted in Figure 2, as a running example. In Figure 2a, we observe a sequential execution, despite T_1 and T_2 accessing different states, limiting throughput. This is, for instance, the case in Ethereum [12] where transactions are executed sequentially. In contrast, Figure 2b illustrates parallel execution, where T_1 , T_2 and T_3 execute in parallel and improve throughput. However, T_4 has a *reads-from* conflict with T_2 on X_3 and therefore must wait for T_2 to commit. Moreover, T_4 has a *write* conflict with T_3 and must either wait for T_3 to commit or employ data structures that allow tracking of both writes performed to X_5 (à la *multi-versioning* [18]).

As we constructively observe, the states accessed by these transactions can be efficiently derived a priori to build *sound*, but possibly incomplete specifications. Depending on how complete the derived specifications are, this information can be used to minimize runtime transactional conflicts and thus avoid unnecessary transaction aborts during parallel execution. However, as the results in this paper show, not only are there rigorous theoretical approaches

to deriving conflict specifications in smart contract ecosystems like Ethereum, fast techniques exist to identify a non-trivial number of transactional conflicts within a block, which provide significant speedup for parallel execution.

4 BTM Execution From Conflict Specifications

In this section, we formally present block transactional memory in the asynchronous shared memory model, detail our BTM algorithms that leverage conflict specification and satisfy preset serializability.

4.1 Model

The model of BTM is presented in the standard asynchronous shared memory model [22].

Transactions. A *transaction* is a sequence of *transactional operations*, reads and writes, performed on a set of virtual machine (VM) states. A BTM $[1, \dots, n]$ *implementation* provides a set of concurrent *processes* with deterministic algorithms that implement reads and writes on accounts using a set of *shared memory locations* accessed by the n transactions with a preset order $T_1 \rightarrow \dots, T_n$. More precisely, for each transaction T_k , a BTM implementation must support the following operations: $read_k(X)$, where X is an object, that returns a value in a domain V or a special value $A_k \notin V$ (*abort*), $write_k(X, v)$, for a value $v \in V$, that returns *ok* or A_k , and $tryC_k$ that returns $C_k \notin V$ (*commit*) or A_k . The transaction T_k completes when any of its operations returns A_k or C_k .

Executions, Histories and Conflicts. A BTM *execution* is a sequence of *events* performed on shared memory states by an interleaving of transactions as prescribed by the implementation. To avoid introducing additional technical machinery that is not strictly necessary to follow the algorithmic exposition in this paper, we do not define executions using configuration semantics as is common in traditional shared memory systems [18, 22], although it would be straightforward to do so.

A BTM *history* is the subsequence of an execution consisting of the invocation and response events of operations of the transactions. For a transaction T_k , we denote all objects accessed by its read and write as $Rset(T_k)$ and $Wset(T_k)$, respectively. We denote all the operations of a transaction T_k as $Dset(T_k)$. The *read set* (resp., the *write set*) of a transaction T_k in an execution E , denoted $Rset_E(T_k)$ (resp. $Wset_E(T_k)$), is the set of objects that T_k attempts to read (resp. write) by issuing a read (resp. write) invocation in E . The *data set* of T_k is $Dset(T_k) = Rset(T_k) \cup Wset(T_k)$. T_k is called *read-only* if $Wset(T_k) = \emptyset$; *write-only* if $Rset(T_k) = \emptyset$ and *updating* if $Wset(T_k) \neq \emptyset$.

We say T_i and T_j *conflict* in an execution E if there exists a common state X in $Dset(T_i)$ and $Dset(T_j)$ such that X is contained within $Wset(T_i)$ or $Wset(T_j)$, or both. Furthermore, we say that T_i, T_j *read-from conflict* if $Wset(T_i) \cap Rset(T_j) \neq \emptyset$ and T_i appears before T_j in the preset order. Note that the definition of read-from conflict, unlike that of a conflict, relies on a preset order existing between the two transactions.

Let H be a sequential history, i.e., no two transactions are concurrent in H . For every operation $read_k(X)$ in H , we define the *latest written value* of X as follows: if T_k contains a $write_k(X, v)$ that precedes $read_k(X)$, then the latest written value of X is the value of the latest such write to X . Otherwise, the latest written value of X is the value of the argument of the latest $write_m(X, v)$ that precedes $read_k(X)$ and belongs to a committed transaction in H . This write is well-defined since H can be assumed to start with an initial transaction

writing to all states. We say that a sequential history S is *legal* if every read of a state returns the *latest written value* of this state in S . It means that sequential history S is legal if all its reads are legal.

Preset Serializability. Given a set of n transactions with a *preset* order $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_n$, we need a deterministic parallel execution protocol that efficiently executes block transactions using the serialization order and always leads to the same state, even when executed sequentially. We formalize this using the definition of *preset serializability*.

► **Definition 1** (Preset serializability). *Let H be a history of a Block-STM $[1, \dots, n]$ implementation M . We say that H is preset serializable if H is equivalent to a legal sequential history S that is $H_1 \cdots H_i \cdot H_{i+1} \cdots H_n$ where H_i is the complete history of transaction T_i . We say that a Block-STM $[1, \dots, n]$ implementation M is preset serializable if every history of M is preset serializable.*

4.2 Algorithm overview

The core idea behind our algorithm is that only independent transactions are executed in parallel, ensuring that no race conditions arise during execution. For any transaction T_k , if there exists a transaction $T_i \notin \text{cSet}(T_k)$ that precedes T_k in the preset serialization order, the execution of T_k is deferred until T_i completes (here $\text{cSet}(T_k)$ denotes the set of transactions that do not conflict with T_k). Additionally, each transaction is followed by validation to ensure that, if two dependent transactions are executed in parallel due to incorrect conflict specification, the one higher in the preset order can abort and re-execute. Thus, the output of the algorithm will be the same as that of a sequential execution.

Performance improvement over parallel execution techniques based on Block-STM is achieved by reducing the number of aborts and re-executions. This is made possible by leveraging the knowledge of the transaction conflicts: a transaction T_k is executed only after ensuring that all preceding transactions in the preset order are in its independence set, $\text{cSet}(T_k)$. This targeted execution strategy minimizes conflicts and improves throughput.

Problem Statement. In this section, we ask the following question: given a set of n transactions $T_1 \rightarrow \dots \rightarrow T_n$ and $\text{cSet}(T_k)$ for all $k \in \{1, \dots, n\}$, what are the most efficient algorithms for implementing Block-STM $[1, \dots, n]$. We implemented two different algorithms, leveraging conflict specification in DAG BTM (dBTM) and optimized BTM (oBTM), based on how the *scheduler* utilizes the transaction conflict specification for efficient parallel execution. As explained in Sections 1 and 2, this mechanism for implementing BTM is applicable directly to the read-write aware models like Sui and Solana in which the BTM is explicitly provided the transactional read-write sets. More importantly, as we demonstrate in Section 6, we can efficiently construct conflict specifications for read-write oblivious models like Ethereum EVM and Aptos MoveVM, even if under-approximate (i.e., incomplete), and still reap the benefits of our proposed methodologies for implementing BTM.

Leveraging Conflict Specification in DAG BTM (dBTM)

This approach utilizes the conflict specifications (assumed to be correct or overapproximated) and the preset order to create a DAG, which serves as a partial order and is provided as input for the scheduler-less execution algorithm. In the DAG, transactions are represented as vertices, whereas conflicts among transactions are denoted as directed edges. The *indegree* field is added with each vertex (transaction) to track dependencies with prior transactions in the preset order; a transaction becomes eligible for execution when its indegree is zero. A

transaction is independent if its $cSet$ contains all preceding transactions in the preset order, resulting in an indegree of zero. During execution, the non-zero indegree transactions wait for the preceding transactions to commit and clear dependencies. At the time of commit, the committing transaction decreases the indegree of all dependent transactions.

We now present the the preset-serializable dBTM (Algorithm 1), which leverages conflict specifications to construct a transaction DAG representing a preset-serializable partial order.

■ **Algorithm 1** dBTM[1, . . . , n]: It is a DAG-based scheduler to execute independent transactions in parallel. The *indegree* field of each transaction vertex in the DAG tracks dependencies on prior transactions. Consider a transaction T_k being executed by a process p_k .

Input: T : list of transactions in the block B_i ; S : pre-state— state before execution of block B_i ; $cSet$: specifications for transactions in T .

Shared Locations: *indegree*[T_k]: the number of transactions in preset order T_i is dependent on; *dependents*[k]: the set of transactions dependent upon T_k ; a version list $\langle v_j \rangle$ for each state X_j .

```

1: Fun readk( $X_j$ ):
2:   if  $X_j \notin Wset(T_k)$  then
3:     // Read the latest version of  $X_j$ 
4:     // created by a  $T_i$  that precedes  $T_k$ .
5:      $[ov_j, i] := read\_lvp(T_k, X_j)$ 
6:      $Rset(T_k) := Rset(T_k) \cup \{X_j, [ov_j, i]\}$ 
7:     return  $ov_j$ 
8:   else
9:     //  $X_j$  is in  $Wset(T_k)$ 
10:     $[ov_j, \perp] := Wset(T_k).locate(X_j)$ 
11:    return  $ov_j$ 
12: Fun readlvp( $T_k, X_j$ ):
13:    $[ov, i] := [0, 0]$ ;
14:   // Read the largest version of  $X_j$ 
15:   // created by a  $T_i$  that precedes  $T_k$ .
16:   forall  $[ov_j, i] \in X_j$  do
17:     if  $k > i$  then
18:        $[ov, i] := [ov_j, i]$ 
19:   return  $[ov_j, i]$ 
20: Fun writek( $X_j, v$ ):
21:    $nv_j := v$ 
22:   if  $X_j \notin Wset(T_k)$  then
23:      $Wset(T_k) := Wset(T_k) \cup \{X_j, [nv_j, k]\}$ 
24:   else
25:     //  $X_j$  is in  $Wset(T_k)$ , update its
26:     // current value to  $v$ .
27:      $Wset(T_k) := Wset(T_k).update(X_j, [nv_j, k])$ 
28:   return ok
29: Fun tryCk():
30:   // Ensure commit order
31:   if indegree[ $k$ ]  $\neq 0$  then
32:     wait until indegree[ $k$ ] = 0;
33:   // Write back to shared memory
34:   forall  $X_j \in Wset(T_k)$  do
35:     Write( $X_j, [nv_j, k]$ )
36:   // Clear dependencies
37:   forall  $T_i \in dependents(i)$  do
38:     indegree[ $i$ ]  $\leftarrow indegree[i] - 1$ ;
39:   return  $C_k$ 
40: Fun gendag( $cSet$ ):
41:   block_size  $\leftarrow size\_of(B_i)$ ;
42:   for  $k \in (0, block\_size - 1)$  do
43:     dependencies[ $k$ ] :=  $\emptyset$ 
44:     dependents[ $k$ ] :=  $\emptyset$ 
45:     indegree[ $k$ ]  $\leftarrow 0$ ;
46:   // Compute complement sets
47:   for  $k \in (0, block\_size - 1)$  do
48:     cSetComp( $T_k$ )  $\leftarrow \{j \mid j < k \text{ and } j \notin cSet(T_k)\}$ ;
49:   // Use cSetComp to build dependencies
50:   for  $k \in (0, block\_size - 1)$  do
51:     foreach  $j \in cSetComp(T_k)$  do
52:       dependencies[ $k$ ].add( $j$ )
53:       dependents[ $j$ ].add( $k$ )
54:   indegree[ $k$ ] :=  $|dependencies[k]|$ 
55:   return (dependents, indegree)

```

Implementation state. The *DAG* is implemented using two primary data structures: 1. *indegree*, where *indegree*[k] denotes the number of preceding transactions in the preset order that transaction T_k depends on; and 2. *dependents*, which maps each T_k to a set of its dependent successors. For each *VM state* X_i , the algorithm maintains a memory location v_i that stores a set of tuples $([v_1, k], [v_2, k'], \dots)$. Each tuple $[v, k]$ represents the value v of X_i and k is the transaction T_k that wrote that value.

Read implementation. The *read_k*(X_j) function returns the value of state X_j as visible to transaction T_k . It first checks whether X_j is already present in $Wset(T_k)$ at line 2. If not, it reads the largest version written by a transaction T_i in shared memory such that

$T_i \rightarrow T_k$, using the helper function $read_lvp()$ defined in lines 9–14. This value, along with its source, is then recorded in $Rset(T_k)$ at line 4. If X_j has already been locally written by T_k , it is returned directly from the $Wset(T_k)$ at line 7. The helper function $read_lvp(T_k, X_j)$ (lines 9–14) is invoked by $read_k$ to identify the latest version of X_j in the multi-version data structure created by a transaction that precedes T_k in the DAG. It iterates over all versions of X_j in shared memory and selects the one with the highest index $i < k$. This mechanism ensures multi-version consistency by enforcing serialization semantics across transactions.

Write implementation. The $write_k(X_j, v)$ function (lines 15 to 21) stages a write to VM state X_j with value v in the local context of T_k . If this is the first write to X_j by T_k , an entry is created in $Wset(T_k)$. If X_j has been written previously, the existing entry is updated with the new value. This operation does not alter shared memory but updates the local write set.

Commit implementation. The $tryC_k()$ function, lines 22 to 29, performs the commit logic for transaction T_k . It first waits until all dependencies are satisfied, that is, all preceding transactions in the DAG on which T_k depends are committed and $indegree[k] := 0$ (line 24). Then updates the shared memory with writes from $Wset(T_k)$ at line 25. At line 27, it clears the dependencies of all successor transactions, specifically, for each transaction $T_i > T_k$ where $T_k \notin cSet(i)$, it decrements $indegree[i]$ by 1 (line 28). Finally, it returns C_k at line 29.

DAG Generation. The $gen_dag(cSet)$ function, lines 30 to 43, constructs a DAG from the specification $cSet$, which maps each transaction T_k to the set of preceding transactions (in preset order) that T_k is independent of. It initializes the per-transaction data structures ($dependencies := \emptyset$, $dependents := \emptyset$, and $indegree \leftarrow 0$) in lines 32–35. Then, computes the complement conflict set $cSetComp(T_k)$ for each transaction $T_k \in T$ (line 36), which includes all preceding transactions of T_k that are not in $cSet(T_k)$. These are the transactions with which T_k conflicts, and hence it must be conservatively dependent on them in the DAG. Using $cSetComp$, the function constructs the dependency edges (lines 38–42). For every $T_k \in T$, the function sets $indegree[k]$ to the cardinality of the dependency set of T_k . Finally, the function returns references to the $dependents$ and $indegree$ structures at line 43, which are used as input to DAG-based scheduler for execution.

Proof of Correctness. We prove preset serializability for the Algorithm 1 by adapting the proof structure from traditional STM literature for the block transactional model [18, 19]. The proof is provided in the extended version on arXiv due to space limitations.

Leveraging Conflict Specification in Optimized BTM (oBTM)

The oBTM algorithm optimizes the Block-STM scheduler by integrating independence information (i.e., $cSet$) associated with transactions, which is assumed to be complete and provided as input. The scheduler is modified to allow transactions that are independent of all preceding transactions in the preset order to execute without validation. A greater number of such independent transactions improves the execution efficiency. This method builds on the existing scheduler of Block-STM [17] for MoveVM and PEVM [28] developed by the RISE chain [27] for EVM. In the absence of conflict specifications, the algorithm defaults to optimistic execution as in PEVM and Block-STM.

The design of oBTM, shown in Algorithm 2, leverages the conflict specifications to enable efficient optimistic execution by minimizing the validation overhead of vanilla PEVM and Block-STM. Consider a transaction T_k being executed by a process p_k .

■ **Algorithm 2** oBTM[1, ..., n]: The approach optimizes the Block-STM scheduler to allow transactions that are independent of all previous transactions to execute without validation. Consider a transaction T_k being executed by a process p_k .

Input: T : list of transactions in the block B_i ; S : pre-state— state before execution of block B_i ; $cSet$: specifications for transactions in T .

Shared Locations: $cSet(T_k)$: set of transactions T_i such that $i < k$ and T_k is independent of T_i .

```

1: Fun  $read_k(X_j)$ :
2:   if  $X_j \notin Wset(T_k)$  then
3:     // Read the largest version of  $X_j$ 
4:     // created by a tx  $T_i$  preceding  $T_k$ .
5:      $[ov_j, i] := read\_lvp(T_k, X_j)$ 
6:     if  $[ov_j, i].is\_estimate$  then
7:       if  $\neg add\_dependency(T_k, T_i)$  then
8:          $\_retry(T_k)$ 
9:       // If  $T_k$  read an estimated value
10:      // written by  $T_i$ 
11:      return  $A_k$ 
12:    $Rset(T_k) := Rset(T_k) \cup \{X_j, [ov_j, i]\}$ 
13:   return  $ov_j$ 
14: else
15:    $[ov_j, \perp] := Wset(T_k).locate(X_j)$ 
16:   return  $ov_j$ 
17: Fun  $write_k(X_j, v)$ :
18:    $ov_j := v$ 
19:   if  $X_j \notin Wset(T_k)$  then
20:      $Wset(T_k) := Wset(T_k) \cup \{X_j, [nv_j, k]\}$ 
21:   else
22:     //  $X_j$  is in  $Wset(T_k)$ , update
23:     // current value to  $v$ .
24:      $Wset(T_k) := Wset(T_k).update(X_j, [nv_j, k])$ 
25:   return  $ok$ 
26: Fun  $tryC_k()$ :
27:   if  $\exists T_i \in cSet(T_k) : i < k$  then
28:     forall  $X_j \in Wset(T_k)$  do
29:       // Update  $Wset(T_k)$  in shared memory
30:        $write(X_j, [nv_j, k])$ 
31:     // Read set validation
32:     if  $\exists X_j \in Rset(T_k) : [ov_j, k] \neq read\_lvp(T_k, X_j)$  then
33:       // On validation failure, versions of
34:       // each  $X_j \in Wset(T_k)$  are marked as
35:       // estimated in shared memory.
36:       forall  $X_j \in Wset(T_k)$  do
37:          $mark\_estimate(nv_j, k)$ 
38:       return  $A_k$ 
39:     return  $C_k$ 
40:   else
41:     //  $T_k$  is an independent transaction.
42:     forall  $X_j \in Wset(T_k)$  do
43:        $write(X_j, [nv_j, k])$ 
44:     // Skip validation for  $T_k$  and commit.
45:     return  $C_k$ 
46: Fun  $add\_dependency(id, blocking\_id)$ :
47:   if  $txn\_status[blocking\_id] == Executed$  then
48:     return  $false$ 
49:   return  $true$ 

```

Implementation state. For each state X_i , the algorithm maintains a memory location v_i storing tuples $([v_1, k], [v_2, k'], \dots)$, where $[v, k]$ represents a value v written to X_i by T_k .

Read implementation. The $read_k(X_j)$ function returns the most recent committed value of X_j visible to transaction T_k . If $X_j \notin Wset(T_k)$, the latest version written by a predecessor T_i is read from shared memory at line 3 using function $read_lvp()$ (as discussed in dBTM algorithm). If the version is an estimate and T_i has not yet committed, $add_dependency$ is invoked to check if T_k must wait (line 5); if so, T_k retries at line 6. Otherwise, the version is added to $Rset(T_k)$ for subsequent validation at line 8. If $X_j \in Wset(T_k)$, meaning it is already in transaction's write set, the value is returned directly from $Wset(T_k)$ at line 11.

Write implementation. The $write_k(X_j, v)$ performs a write that assigns a value v to state X_j by transaction T_k . It first checks whether $X_j \notin Wset(T_k)$, if so, a new entry $[v, k]$ is inserted into $Wset(T_k)$ (line 16); otherwise, the existing entry is updated with the new value v (line 18). This operation does not update in shared memory; instead, it maintains a local view of writes until commit (line 19).

Commit implementation. The $tryC_k()$ function attempts to commit transaction T_k . If T_k depends on any predecessor transaction(s) in preset, it first updates its $Wset(T_k)$ to shared memory (a multi-version data structure) at line 23 and then validates its read set $Rset(T_k)$

against the latest values committed to shared memory at line 24. If validation fails, the versions written by T_k in the multi-version structure are marked as estimated at line 26, and T_k aborts by returning A_k at line 27. If validation succeeds, or if T_k is an independent transaction (its $Wset(T_k)$ is committed to shared memory and validation is skipped), the function commits T_k by returning C_k at line 28 and line 32, respectively.

5 Implementation and Evaluation of BTM from Conflict Specifications

In this section, we analyze the execution latency and throughput of dBTM and oBTM, compared to baseline sequential execution and state-of-the-art parallel execution techniques.

5.1 Implementation

EVM Implementation. We compare dBTM and oBTM with PEVM [28], a version of Block-STM [17] for the EVM developed by the RISE chain [27]. Experiments are conducted on REVM [8] version 12.1.0, an EVM implementation written in Rust. We used the core data structures provided by PEVM version 0.1.0 (commit hash f0bdb21) and implemented our algorithms on top of them. Specifically, we used a hash map to store the cSet, which allowed us to store both the list of dependent transactions and the indegree of each transaction in the DAG. A custom parallel-queue is implemented to manage and enable efficient parallel execution across multiple threads.

We analyzed the performance on both synthetic and historical blocks from Ethereum’s mainnet. In the EVM, each transaction pays a gas fee to the *Coinbase* account, which belongs to the block proposer. As a result, every transaction updates the Coinbase account, leading to a 100% write-write conflict with all preceding transactions. To address this, we defer the fee transfer; that is, the gas fees are locally collected per transaction and credited to the Coinbase account only at the end, after all transactions in the block have been executed.

MoveVM Implementation. We analyzed the performance of our dBTM and oBTM approaches against the baseline Block-STM [17], which runs DiemVM (an earlier version of MoveVM) using the test setup from [7]. The evaluation includes two primary workloads: (i) peer-to-peer (P2P) transfers, as provided in the original test setup [7], and (ii) synthetic workloads (batch transfers and a generic mixed workload) we designed to mimic diverse execution patterns. Due to space constraints and for presentation succinctness, we defer the detailed analysis of BTM on MoveVM to the arXiv version.

Experimental Setup. We ran experiments on a single-socket AMD machine consisting of 32 cores (64 vCPUs), 128 GB of RAM, and running Ubuntu 18.04 with 100 GB of SSD. The experiments are carried out in an execution setting of in-memory using REVM and MoveVM, and abstract away the latency of persistent storage access. We ran experiments 52 times, with the first 2 runs designated as warm-up. Each data point in the plots represents an average throughput (tps) and latency (ms), where each execution is repeated 50 times.

5.2 EVM Analysis

Figure 3 to Figure 6 show the relationship between varying threads or varying conflicts (α -tail distribution factor) within the block, represented on the X-axis, on two important performance metrics: *throughput (tps)*, represented by histograms on the primary Y-axis (Y1), and *execution latency (ms)*, as line graphs on the secondary Y-axis (Y2).

Construction and Testing of Synthetic Workloads

We conducted experiments on synthetic workloads to analyze throughput and latency under various execution conditions. A synthetic workload is a computational workload that does not depend on user input but instead simulates transaction types and concurrency patterns.

We analyzed Ethereum’s Mainnet blocks to generate synthetic workloads and found that certain accounts (or smart contracts) are predominantly accessed, indicating that historical blocks exhibit a tail distribution in address access frequency. Based on this observation, we designed synthetic workloads that reflect varying characteristics of transaction conflicts in blocks. Following the tail distribution, we evaluated performance across synthetic workloads, including worst-case and best-case scenarios, enabling a comprehensive assessment of system behavior under varying levels of conflict within a block. For this, we added another metric on the X-axis, the parameter α , which represents the heaviness of the tail in the Pareto distribution [24]. The value of α determines the degree of conflict, as shown in subfigure (b); smaller values of α correspond to fewer conflicts.

1. ERC-20 Workload (W_{erc20}): This workload simulates multiple ERC-20 smart contracts, with each contract representing a distinct cluster. Transfer transactions occur among addresses within each cluster, each of which has its own ERC-20 token. The number of transactions per cluster is determined using a Pareto distribution [24], while the sender and receiver addresses are selected uniformly within each cluster. This approach captures the distribution of transaction loads, resulting in certain clusters receiving higher transactions. The conflict specifications are derived using the sender, receiver, and contract addresses accessed by transactions in the block. For testing, we generate 10k ERC-20 transfers across 10k contract addresses (clusters), with each cluster mapped to a unique externally owned account (EOA) address. Since each contract has a single EOA, transactions are effectively self-transfers. However, an EOA may initiate multiple transactions within a cluster, potentially causing conflicts, with a tail factor of $\alpha = 0.1$.

As shown in Figure 3a, increasing thread count improves performance. In fact, dBTM achieves peak throughput at 24 threads with 411k tps, compared to 178k tps and 171k tps for oBTM and PEVM, with an average throughput of 299k, 157k and 143k tps, respectively. In particular, dBTM achieves a $4\times$ speedup over sequential and a speedup of $2.4\times$ over PEVM. Latency trends show that optimistic approaches initially reduce execution time, but later increase it due to higher abort rates. In contrast, dBTM minimizes aborts by leveraging conflict specifications, reducing both abort and (re-)validation costs. Similarly, in Figure 3b, with 16 threads, increasing block conflicts leads to a steady performance decline, which is more pronounced in optimistic approaches due to higher abort rates and increased waiting in dBTM. At $\alpha = 5$ (maximum conflicts), all approaches incur a higher overhead and perform worse than sequential execution.

2. Mixed Workload (W_m): This workload combines 50% native ETH transfers and 50% ERC-20 transfers, ensuring a balanced representation of both types of transactions. For ETH transfers, sender and receiver addresses are selected using a Pareto distribution [24]. The number of ERC-20 contracts is set to 25% of the block size, with workload generation similar to W_{erc20} . As a result, 2.5k contracts (clusters) comprise a block of 10k transactions.

As shown in Figure 4, throughput and latency remain consistent with W_{erc20} workload. However, due to 50% ETH transfers (microtransactions) in the block, all approaches show a noticeable increase in throughput. The average tps is 356k for dBTM, 238k for oBTM, 229k for PEVM, and 157k for sequential. This translates to speedups of $2.27\times$, $1.52\times$, and $1.46\times$ for dBTM, oBTM, and PEVM, respectively, over sequential execution. The maximum tps of dBTM increased from 411k to 454k.

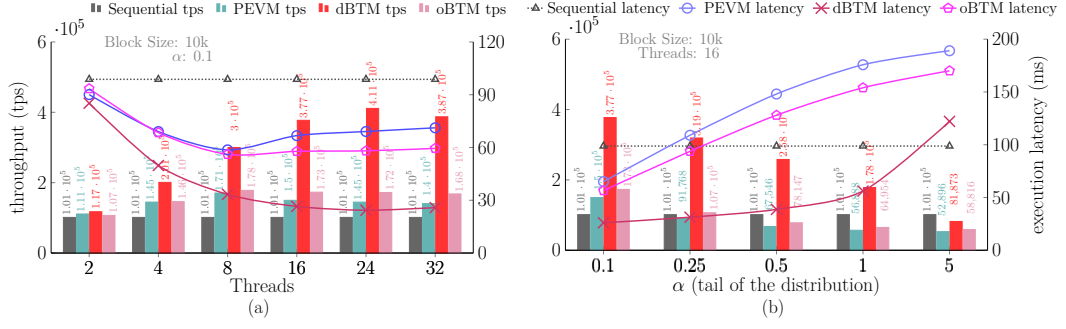


Figure 3 ERC-20 transfer workload (W_{erc20}).

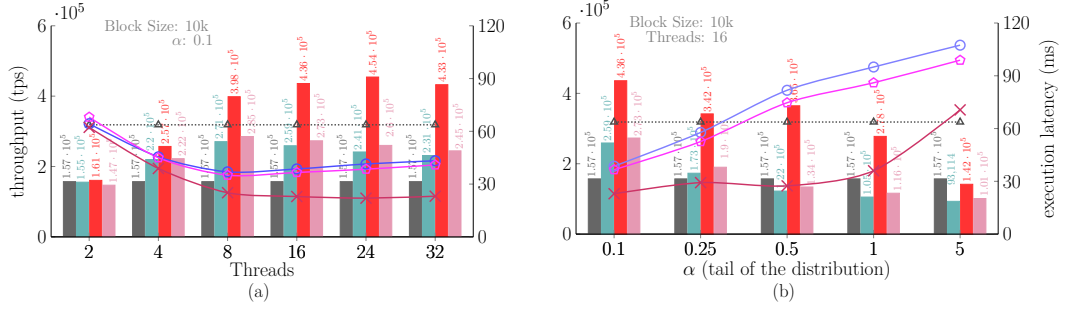


Figure 4 Mixed workload (W_m).

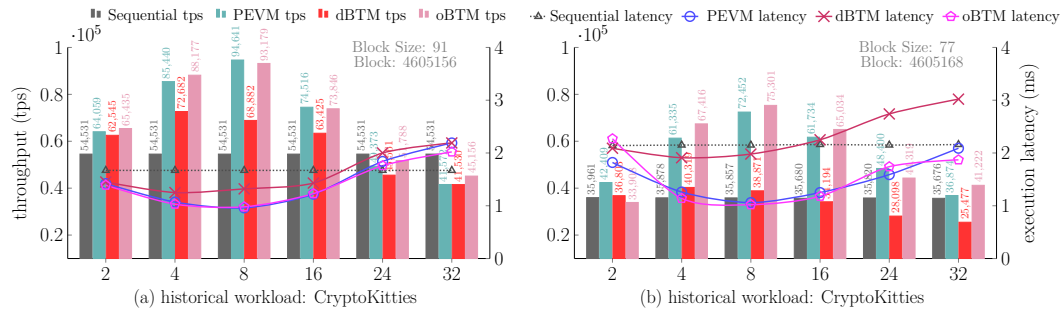
Testing on Real-World Ethereum Transactions

We selected blocks from different historical periods based on major events that may have impacted Ethereum’s concurrency and network congestion. Each of these blocks allows us to analyze performance in different historical periods, providing insight into how major events, such as popular dApp launches and significant protocol upgrades, affect transaction throughput and network latency. Consequently, this also helps us understand the limitations of parallel execution approaches under different network conditions, such as the gradual increase in conflicts in the block and the change in resource requirements over time.

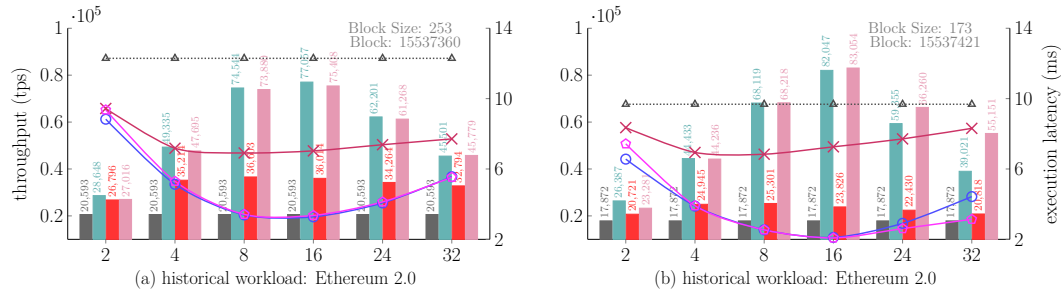
We trace the states accessed by transactions within a historical block using **callTracer** and **prestateTracer** APIs [9], which provide a complete view of the block’s pre-state, the state required for executing the current block. To derive conflict specifications, the pre-state file is parsed to identify all EOAs and smart contract addresses accessed within the block. These addresses form the basis for constructing the conflict specification and serve as the “ground truth”. Note that to derive specifications for a transaction T_k , all transactions (their pre-state EOAs and smart contract addresses) preceding it in preset order are considered.

Note that real historical Ethereum blocks contain significantly fewer transactions than to our synthetic tests. To ensure a representative evaluation, we selected larger blocks from different historical periods. Specifically, we analyzed two blocks each from the CryptoKitties contract deployment and Ethereum 2.0 merge periods. We also analyzed two recent blocks with specific characteristics to demonstrate the worst-case and the best-case parallelism.

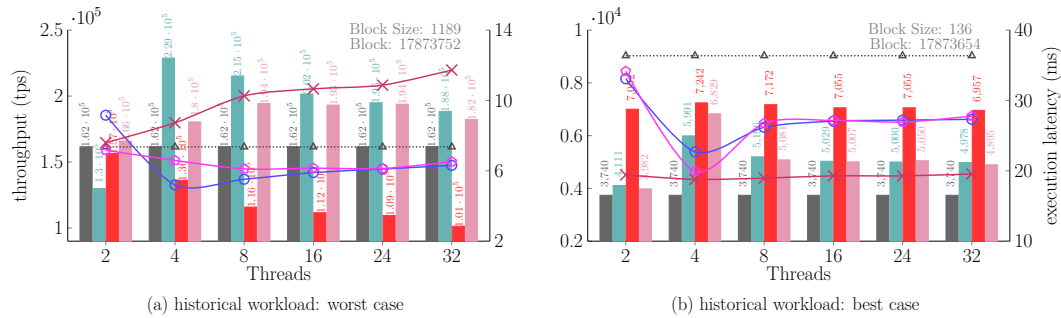
1. CryptoKitties Contract Deployment (W_{ck}): The famous CryptoKitties [10] contract was deployed on block 4605167, after which an unexpected spike in transactions caused Ethereum to experience high congestion. We analyze block 4605156, which occurred before the contract deployment, and block 4605168, which took place after the deployment.



■ Figure 5 CryptoKitties historical period (W_{ck}).



■ Figure 6 Ethereum 2.0 merge historical period (W_{e2}).



■ Figure 7 EVM: Historical blocks of worst and best-case (W_{wb}).

Figure 5a demonstrates a significant speedup in parallel execution, while Figure 5b highlights congestion in the block following contract deployment. In dBTM, conflicts are over-approximated due to sender and contract address-level conflict specifications, as real blocks lack read-write and contract storage access-level granularity. This leads to a performance drop compared to synthetic tests. With complete conflict specifications and larger blocks, dBTM could achieve superior performance, as observed in synthetic tests. That said throughput increases with additional threads, peaking at 8 threads, indicating the optimal point for speedup. Notably, oBTM outperforms other approaches in this workload both before and after contract deployment. The average tps for oBTM, PEVM, and dBTM drops from 69k, 68k, and 59k, before contract deployment to 54k, 53k, and 33k after deployment. Additionally, despite smaller blocks in the post-contract deployment, execution time increases for both sequential and other executors, highlighting the impact of congestion on overall performance.

2. Ethereum 2.0 Merge (W_{e2}): The merge [13] took place in block number 15537393, which changed Ethereum’s consensus to proof-of-stake along with other protocol-level changes and had an impact on transaction processing, block validation and network traffic in general. We analyze block 15537360 before the merge and 15537421 after the merge. Compared to the W_{ck} historical period, the network had experienced an increase in block size and user activity. As shown in Figure 6, the average throughput for oBTM, PEVM, and dBTM before the merge was 55k ($2.68\times$ speedup over sequential), 56k ($2.73\times$ speedup), and 33k ($1.63\times$ speedup), respectively. Post-merge, these values changed to 56k ($3.17\times$ speedup), 53k ($2.98\times$ speedup), and 23k ($1.29\times$ speedup). Notably, the increased speedup over sequential post-merge shows the growing parallel execution potential to improve throughput.

3. Worst-case and Best-case Blocks (W_{wb}): To analyze performance extremes, we selected Ethereum blocks exhibiting different conflict profiles. The worst-case block, 17873752, represents maximal conflicts with densely interconnected transactions and consists of 1189 transactions, of which 1129 are ETH transfers, 21 are ERC-20 transfers, and 39 are other contract transactions. We observed that most ETH transfers interact with the same address, resulting in a longer dependency chain of length 1104. As a result, the block becomes inherently sequential, and hence a good choice for worst-case analysis. In contrast, the best-case block, 17873654, consists predominantly of independent transactions. There are a total of 136 transactions, including 41 ETH transfers, 18 ERC-20 transfers, and 77 other contract transactions. We chose this block because it has a significant number of internal transactions (435) and the length of the longest conflict dependency chain is 50.

As shown in Figure 7a, in the worst-case, for dBTM the tps remains consistently low, around 100k-137k across all thread counts and 166k-195k for oBTM. This reflects minimal improvement over the sequential tps of 162k, highlighting the limitation of parallel execution under blocks with high access-level dependencies. While PEVM at 4 threads shows $1.41\times$ performance gains over sequential, increasing the threads beyond introduces aborts and re-executions, resulting in overhead. However, the performance gain could be due to the lazy update of ETH transfers, which defers balance changes to EOAs until a later stage at commit time, as well as to independent contract transactions that are computationally intensive compared to ETH transfers—microtransactions with minimal computational logic. Due to the overhead of DAG construction, dBTM shows a noticeable performance degradation. Moreover, 94.95% of the transactions in the block are ETH transfers resulting in smaller execution latency despite the large block size. The average execution latency for PEVM is 6.36 ms, oBTM is 6.45 ms, that are not significantly less compared to 7.36 ms of sequential execution, while dBTM incurs slightly higher latency 9.97 ms.

In contrast, the best-case block in Figure 7b shows that oBTM achieves throughput in the range of 3.9k-6.8k tps, while dBTM ranges from 6.9k-7.2k tps, both significantly outperforming 3.74k tps of the sequential execution. The average execution latency across threads is 27.1 ms for oBTM and 19.2 ms for dBTM, compared to 36.36 ms for sequential and 27.25 ms for PEVM. In terms of average speedup, oBTM achieves a $1.34\times$ improvement over sequential and performs on par with PEVM, while dBTM shows superior efficiency, with an average speedup of $1.89\times$ over sequential and a $1.42\times$ speedup over PEVM. Since, this block contains a higher proportion of smart contract transactions with numerous internal calls, making it computationally intensive for the VM. Consequently, minimizing aborts and re-executions yields tangible performance benefits. Notably, dBTM maintains its competitive edge due to its reduced abort and re-execution overheads.

Observe that from one historical period to another (even with just these blocks), despite the block size increasing, the overall throughput has decreased, highlighting the impact of congestion and resource requirements on overall performance.

Summary. *Our EVM analysis highlights the strong performance of dBTM and oBTM, demonstrating their potential to achieve a near-theoretical maximum parallel execution by leveraging conflict specifications. With more accurate conflict specifications in historical workloads, we can expect to achieve even higher speedups over sequential and PEVM. However, in the absence of conflict specifications, oBTM is expected to perform the same as the optimistic execution of PEVM. On average, across synthetic workloads, dBTM achieves a maximum of $4\times$ speedup over sequential and $3.82\times$ over PEVM, while oBTM achieves $1.8\times$ and $1.2\times$ speedups, respectively. In contrast, across historical blocks, dBTM achieves an average speedup of $1.22\times$ over sequential and a maximum of $1.7\times$ over PEVM, while oBTM achieves an average of $1.83\times$ over sequential and a maximum of $1.41\times$ over PEVM.*

6 Block Transactional Memory for EVM and MoveVM

Having established the empirical advantages of BTM's leveraging conflict specifications, we now present the design of a *full conflict analyzer* integrated implementation of the EVM parallel execution. We demonstrate with the conflict analyzer that it is possible to algorithmically implement *sound* conflict specifications for real Ethereum blocks (cf. Section 6.1). The derived specifications, though might be *incomplete*, i.e., they may not identify conflict relations for some class of transactions, can be generated efficiently and thus prove consequential for maximizing transaction execution throughput (as established in Section 4.2). Additionally, we present results for a workload adaptive execution leveraging the conflict specifications for EVM in Section 6.3. We then present a limited set of performance numbers for conflict analyzer integrated BTM on the MoveVM in Section 6.4. The BTM implementation for MoveVM leverages a custom-built conflict analyzer that is more conservative than the EVM conflict analyzer because the Aptos Move memory model is quite different from EVM posing several challenges for a general specification derivation scheme. The detailed version of the MoveVM implementation and analysis is delegated to the online arXiv version.

6.1 Conflict Specifications for EVM

Preliminaries. We now detail some technical preliminaries needed to describe the approach for building conflict specifications in EVM.

► **Definition 2.** *Ethereum blockchain is made up of a set of **accounts** $A = A_U \cup A_C$, where A_U is the set of EOAs and A_C is the set of contract accounts. Each account $a \in A$ is associated with some data $a.\text{balance}$ such that $a.\text{balance} \in a.\text{data}$ where $a.\text{balance}$ is the ETH balance. For an EOA $a \in A_U$, $a.\text{data} = \{a.\text{balance}\}$ whereas a contract account usually has fields associated in $a.\text{data}$.*

► **Definition 3.** *An Ethereum transaction is specified by a tuple $(\text{origin}, \text{dest}, \text{value}, \text{calldata})$ where $\text{origin} \in A_U$ is the initiator of the transaction, $\text{dest} \in A$ is the transaction recipient, $\text{value} \in \mathbb{N}$ is the number of wei sent by origin to dest , calldata is the accompanying data. An Ethereum transaction T is a simple ETH payment if $T.\text{dest} \in A_U$ and is a contract transaction otherwise. For the read/write sets it always holds that $Rset(T), Wset(T) \subseteq \cup_{a \in A} a.\text{data}$.*

► **Note 4.** A real-world Ethereum transaction contains more elements, but for the purpose of overapproximating read/write sets here, such tuple is a sufficient enough characterization.

► **Definition 5.** For a contract transaction T , the **signature** of T denoted by $T.sig$ is the first four bytes of $T.calldata$ that specify the function to call in the contract ($T.dest$).

► **Example 6.** Consider a situation where an EOA $a \in A_U$ initiates a transaction T_1 calling the function `receiveMessage` of a *Contract* b (cf. Listing 1) so the corresponding tuple is: $(origin : a, dest : b, value : 1, calldata : \{signature : receiveMessage, arg : "hello"\})$. We can see here, that always $Rset(T_1) = \{a.balance, b.balance, b.shouldAccept\}$. Notice that whenever $T.value > 0$ we do have $T.origin.balance, T.dest.balance \in Rset(T)$, $Wset(T)$ as the value needs to be deducted from $T.origin$ and added to $T.dest$. But either $Wset(T_1) = \emptyset$ or $Wset(T_2) = \{b.message\}$ depending on the value of $b.shouldAccept$, because the message field is updated only if the `shouldAccept` value is true.

Now, let another EOA account c initiate a transaction T_2 with the corresponding tuple $(origin : c, dest : b, value : 0, calldata : \{signature : setShouldAccept, arg : false\})$. We have $Rset(T_2) = \emptyset$, $Wset(T_2) = \{b.shouldAccept\}$. We can see from the read and write sets characterized here that T_1, T_2 can never have a read-from conflict if T_1 precedes T_2 (regardless of the control flow in T_1) in the preset order. But if T_2 precedes T_1 , then they will have read-from conflict since T_2 writes to $b.shouldAccept$ and T_1 reads from it.

Approach. The EVM conflict analyzer is just labeling simple payments as *SimplePayment*, the designated ERC-20 functions with their respective signature and everything else as *exitsContract* in line 8. After the labeling, check if a transaction T_j can have read-from conflicts with some $T_i, i < j$, we consider multiple scenarios. If T_j is labeled *exitsContract*, we simply assume that they do have read-from conflict (line 18). If both T_i, T_j are labeled as one of the special ERC-20 transactions and interact with the same contract, it just checks to see if the sender and accounts in the arguments of the functions called for T_j are disjoint from those of T_i (line 25). This is enough as an ERC-20 contract only affects some mapping (*balances, allowances*) in the contract storage with the keys affected being either the transaction arguments or the initiator ($T.origin$).

In all other cases, it is enough to check that the origin and destination of T_i is disjoint from those of T_j (line 30). Since both ERC-20 transactions and simple payments only affect the source and target of the transactions, this is enough to guarantee independence.

We now go over the Algorithm 3 for transactions interacting with *Token contract* and *Wallet contract*, as shown in Listing 1. Notice that every transaction that calls a function other than *Token.transfer* in these two contracts will be labeled as *exitsContract* and thus be assumed to conflict with everything else. Therefore, just consider an account a_1 initiating a transaction T_1 invoking *Token.transfer* with address a_2 and another account b_1 initiating T_2 invoking *Token.transfer* with address a_2 . Here, both functions are labeled *transfer*. Now, to check whether they are independent, one needs to check whether all a_1, a_2, b_1, b_2 are distinct. Now, let T_3 be a simple payment of c_1 to c_3 . Of course T_3 is labeled as *SimplePayment*. Similarly, one only needs to check the distinctness of c_1, c_2 with respect to each of $a_1, Token$ and $b_1, Token$ to see if any conflict arises.

■ **Listing 1** Smart Contracts.

```

1  contract b {
2      string message = "";
3      bool shouldAccept = true;
4      function receiveMessage(String message) external payable{
5          if(shouldAccept){
6              this.message = message;
7          }
8      }
9      function setShouldAccept(bool shouldAccept) external {
10         this.shouldAccept = shouldAccept;
11     }
12 }
13 contract Token {
14     mapping(address => uint) balances;
15     address priceOracle;
16     function transfer(address target, uint amount) external {
17         require(tokenBalances[msg.sender]>=amount);
18         balances[msg.sender] -= amount;
19         balances[target] += amount;
20     }
21     function turnEtherToToken() external payable{
22         (, bytes memory data) = priceOracle.staticcall(abi.encodeWithSignature("tokenPrice()"));
23         uint tokenPrice = bytesToInt(data);
24         balances[msg.sender] = balances[msg.sender] + (msg.value/tokenPrice);
25     }
26 }
27 contract Wallet {
28     mapping(address => uint) balances;
29     function addToWallet() external payable {
30         balances[msg.sender] += msg.value;
31     }
32     function withdraw(uint amount) external {
33         require(balances[msg.sender]>=amount);
34         balances[msg.sender] -= amount;
35         msg.sender.transfer(amount);
36     }
37 }

```

■ **Algorithm 3** Output conflict specifications, given as input to Algorithms 1 and 2.

Input: A transaction T_j in a preset order $T_1 \rightarrow \dots \rightarrow T_n$. **Output:**
 $S \subseteq \{T_i | i < j, Rset(T_j) \cap Wset(T_i) = \emptyset\}$.

<pre> 1: Fun isERC-20 (label): 2: return label $\in \{transfer, transferFrom, approve\}$ 3: Fun erc20Accounts (T): 4: // Get the affected accounts in an ERC-20 5: transaction, for instance, a transferFrom() 6: involves from, to, and the sender address. 7: Two ERC-20 transfers are independent if 8: their affected account sets are disjoint. 9: if label $\in \{transfer, approve\}$ then 10: return {T.origin, T.calldata.to} 11: else if label = transferFrom then 12: return {T.origin, T.calldata.from, T.calldata.to} 13: Fun getLabel (T): 14: if T.dest $\in A_V$ then 15: return SimplePayment. 16: sig $\leftarrow T.sig$ 17: if sig $\in \{transfer, transferFrom, approve\}$ in 18: ERC-20 then 19: return T.sig 20: else 21: return exitsContract </pre>	<pre> 16: Fun findCSet(T_j): 17: label_j \leftarrow getLabel(T_j) 18: if label_j = exitsContract then 19: return \emptyset 20: Output $\leftarrow \emptyset$ 21: forall 1 $\leq i < j$ do 22: label_i \leftarrow getLabel(T_i). 23: if T_i.origin = T_j.origin or label_i = 24: exitsContract then 25: continue; 26: if T_i.dest = T_j.dest and isERC-20 (label_i) 27: and isERC-20 (label_j) then 28: affectedAccounts_i \leftarrow erc20Accounts (T_i) 29: affectedAccounts_j \leftarrow erc20Accounts (T_j) 30: if affectedAccounts_i \cap 31: affectedAccounts_j = \emptyset then 32: Output \leftarrow Output $\cup \{T_i\}$ 33: else if {T_i.origin, T_j.origin, T_i.dest, T_j.dest} = 34: 4 then 35: Output \leftarrow Output $\cup \{T_i\}$ 36: return Output </pre>
--	---

► **Proposition 7.**

- Whether T is a simple payment or calls a designated ERC-20 function, only writes to or reads from accounts $T.origin, T.dest$.
- If T calls an ERC-20 function in transfer, transferFrom, approve, then only $T.dest$ storage is affected. Assuming that the contract adheres to ERC-20 specifications, as long as $T.origin$ and accounts in $T.data$ are disjoint from another T' calling these specific

functions in $T.\text{dest}$, then different parts of the same mappings in $T.\text{dest}$ are modified. Also, if the contract is implemented as a proxy, we assume that only the proxy contract can call the parent (no EOA can initiate a transaction directly with the parent contract).

► **Theorem 8.** *If the Algorithm 3 outputs an independence set \mathcal{T} for a transaction T_j , then for each $T \in \mathcal{T}$ in the preset order before T_j , $Rset(T_j) \cap Wset(T) = \emptyset$.*

Proof. Let $T \in \mathcal{T}$ in the preset order be before T_j . First note that by the algorithm, if T label is *exitsContract*, then it is overapproximated to have read-from conflict with every transaction before in the preset order. So we can assume T is either a simple payment or an ERC-20 transaction. Then (according to the definition of simple payment and Proposition 7 and line 8) T only modifies/reads from $\{T.\text{origin}, T.\text{dest}\}$ which is guaranteed to be disjoint from $\{T_j.\text{origin}, T_j.\text{dest}\}$ by line 30. Therefore, by line 8 and Proposition 7, T_j only modifies $\{T_j.\text{origin}, T_j.\text{dest}\}$ which is disjoint from the read set of T implying that T does not have a read-from conflict with T_j . ◀

6.2 Conflict Analyzer Integration with EVM for Parallel Execution

This section presents the analysis of *full* conflict analyzer integrated with our BTM algorithm for the EVM. We evaluated the performance of the proposed BTM algorithm on the EVM, leveraging conflict specification provided by the conflict analyzer.

Implementation Details. We introduced several implementation-level modifications to the oBTM, resulting in an extended version that we refer to as iBTM. It uses analyzer-provided dependency information (beyond the sender, receiver, and contract addresses used in oBTM) to reduce aborts, and applies the lazy update optimization (of PEVM) for native ETH transfers, while other transactions remain optimistically executed. We choose to perform tests with iBTM, as it required minimal changes to existing implementation and demonstrates the potential benefits of the conflict analyzer when integrated with the optimistic execution of oBTM. Note that for this analysis, we re-run the experiments in our integrated conflict analyzer setup, and the numbers in Section 5.2 can be treated as the ground truth. The entire implementation built on top of PEVM which is 8097 lines of code with 4454 lines added over the original PEVM implementation.

The conflict analyzer is implemented in approximately 700 lines of Rust code. It uses HashMap that maps each accessed address to a vector of transactions that have interacted with it. Transactions labeled as *exitsContract* are pruned in the first pass to reduce the number of pairs evaluated for independence as a transaction labeled *exitsContract* is assumed to conflict with all others. After that if the number of transactions remaining after that is high enough (a manually set parameter), parallelization is done across 16 threads. In addition to the implemented version of the conflict analyzer detailed in Section 6.1, we additionally implemented another *strong* version of the conflict analyzer. Both of these are automated completely and proved to be sound (but not complete) if correctly implemented (cf. Theorem 8). The implemented analyzer main limitation is its inability to deal with non ERC-20 contracts, marking every transaction to such contracts as dependent on all other transactions. Strong analyzer main limitation is that when a contract function calls another contract C not known statically (that is, the address of the C is determined by the input data of the transaction) one cannot know what sequence of blockchain state manipulations after that and thus such transactions are marked dependent on everything else. In the case that C and its payload are known statically, it is rather easy to add more rigorous specifications. There exist inputs which can make the analyzer derive no independence specification making

■ **Table 1** Conflict analyzer integrated with iBTM for parallel execution in EVM; metrics from conflict analysis phase, block composition, and execution performance across different approaches.

	Metric	Ethereum Historical Blocks				Synthetic 10k Blocks	
		4605168 (CryptoKitties)	15537421 (Ethereum 2.0)	17873752 (Worst-case)	17873654 (Best-case)	ERC-20	Mix
Conflict Analysis Phase	Conflict Generation Time	18.314 μ s	39.00 μ s	106.79 μ s	30.09 μ s	6.06 ms	3.30 ms
	Specification Size	627	943	50706	1691	49658226	49824316
Block Stats	Block Size	77	173	1189	136	10000	10000
	ETH Transactions	32	33	1129	41	0	5000
	ERC-20 Transactions	9	11	21	18	10000	5000
	Other Contract Transactions	36	129	39	77	0	0
	Dependency Count	10	7	4	0	5683	3215
Average Execution Time (16 threads)	Sequential	2.107 ms	15.60 ms	7.21 ms	36.63 ms	114.43 ms	72.12 ms
	PEVM	1.32 ms	6.93 ms	6.03 ms	27.52 ms	96.75 ms	53.60 ms
	oBTM	1.24 ms	6.88 ms	5.99 ms	27.68 ms	88.67 ms	54.14 ms
	iBTM	1.23 ms	6.43 ms	5.43 ms	26.85 ms	79.42 ms	39.50 ms

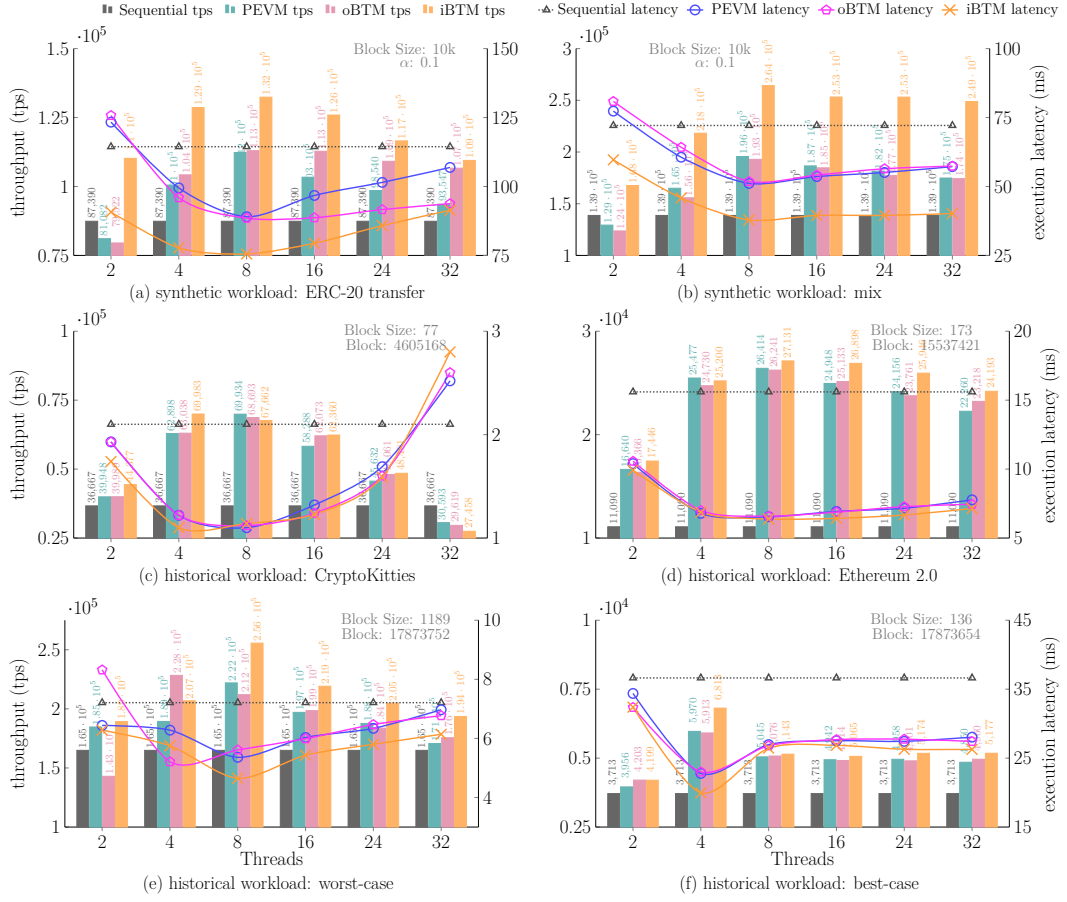
the execution as slow as sequential. To be more precise, we analyze the ERC-20 contracts by their function signature, so a malicious contract not adhering to ERC-20 specifications may result in an unsound analysis, however, one can regularly inspect deployed contracts either manually or automatically to adhere to specifications and analyze only those marked to adhere to specifications (which can be done by current tools [14, 23]).

Testing on Synthetic Workloads. Table 1 presents comprehensive details of integrated experiments, including conflict analysis and execution latency on historical Ethereum and large synthetic blocks. The reported metrics capture statistics from the conflict analysis phase, including the time required to generate conflict specifications and their size, as well as block characteristics such as block size, transaction type distribution, and conflict dependency counts for block. Also report the execution time under three configurations: sequential, PEVM and iBTM (execution with 16 threads). The goal is to evaluate computational cost and execution efficiency in both historical and large synthetic blocks.

We choose two synthetic workloads from Ethereum analysis Section 5.2: ERC-20 and Mix workloads, with a block size of 10k transactions and a fixed conflict distribution factor $\alpha = 0.1$, while varying the thread count on X-axis. Figure 8 shows the advantage of the conflict analyzer in saving the abort and re-execution cost for optimistic execution. As the thread count increases, in both workloads, ERC-20 transfer (Figure 8a) and Mix (Figure 8b) a noticeable execution latency improvement can be seen for iBTM over PEVM, highlighting the benefits of leveraging conflict specifications for execution. The benefit is especially pronounced for the ERC-20 workload, which contains more uniform transaction patterns.

Testing on Historical Ethereum Blocks. Figures 8c–f illustrates the performance on Ethereum blocks. We select one block each from the CryptoKitties (same block as in Figure 5b) and the Ethereum 2.0 workloads (same block as in Figure 6b). We also analyzed performance on the worst and best-case blocks of the workload in Figure 7. As shown in Figure 8, with our conflict analyzer, read-write conflict detection improves the performance of iBTM, which can be clearly observed by comparing the relative performance of PEVM with iBTM. The performance of iBTM improves in almost all cases compared to PEVM and oBTM with a noticeable difference, and is expected to improve further with a fully optimized and mature conflict analyzer. This performance could be due to the reduced number of aborts and re-executions (re-validations) in the optimistic execution.

To conclude, the proposed iBTM achieves up to a $1.75\times$ speedup over sequential execution and a $1.33\times$ speedup over PEVM, with an average speedup of $1.24\times$ over PEVM in synthetic workloads. In historical workloads, it achieves a maximum speedup of $2.43\times$ and $1.14\times$ over sequential and PEVM, also consistently outperforms PEVM across all evaluated scenarios.



■ **Figure 8** Conflict analyzer integrated with EVM: synthetic and historical workloads.

6.3 Adaptive Implementation Based on Conflict Threshold

The advantage of the conflict specification is that it allows us to determine how many pairwise conflicts exist in a block of transactions. We demonstrate how we can leverage this for an adaptive implementation that can also handle high-conflict workloads (Figure 9a). Specifically, by computing a *conflict threshold* for each block with minimal overhead, we deterministically fall back to sequential execution in the case of high conflicts.

Figure 9 demonstrate the advantage of an adaptive technique that dynamically selects the most suitable execution path based on the characteristics of the workload. When $\alpha = 0.1$, the iBTM outperforms all other approaches, leading the adaptive technique to choose it for execution. In contrast, when $\alpha = 5$, the workload becomes highly sequential, and all parallel approaches perform worse than sequential execution. The adaptive approach could fall back to sequential execution; however, there could be small overhead with selecting the optimal approach; this is offset by significant execution time savings in highly conflicting workloads. The current adaptive mechanism is implemented and evaluated for two execution paths, sequential or iBTM, and considers the conflict threshold and block size for decision making. Exploring additional metrics beyond the conflict threshold and block size, such as available compute, expected block computation cost, and past n block execution statistics, as well as different adaptive parallel execution paths (e.g., dBTM, iBTM, PEVM, sequential), remains ongoing work and may constitute standalone future work.

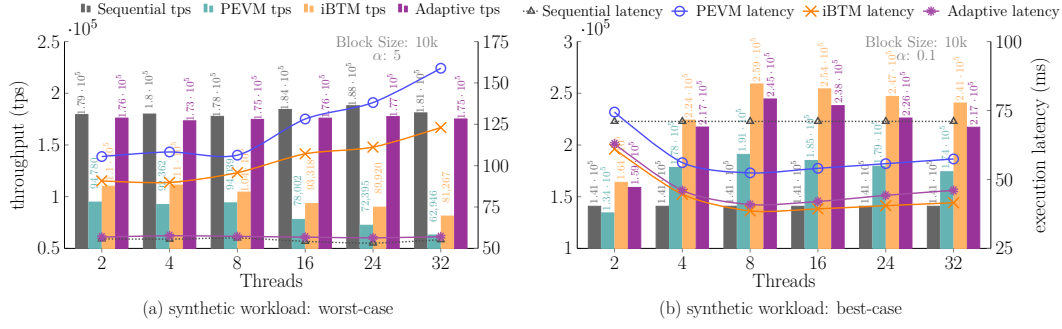


Figure 9 Adaptive on EVM synthetic workloads.

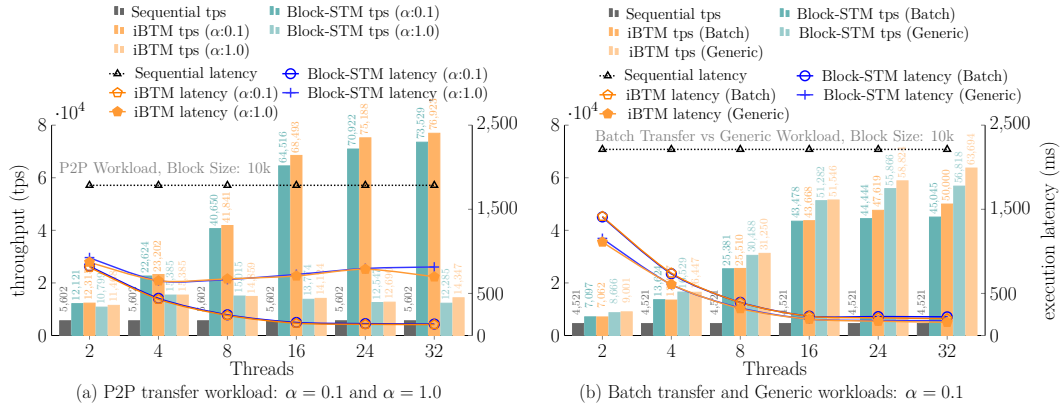


Figure 10 Conflict analyzer integrated with MoveVM: P2P, batch transfer and generic workloads.

6.4 Conflict Analyzer Integration with MoveVM for Parallel Execution

Figure 10 presents the analysis of iBTM for MoveVM on synthetic workloads: P2P, batch, and a generic transfers that combine P2P and batch transfer in equal proportion. Figure 10 shows throughput and latency results for MoveVM across different synthetic workloads leveraging specifications generated by the conflict analyzer. We compared the performance of iBTM with the baseline sequential and the state-of-the-art Block-STM parallel execution while varying the thread count from 2-32, with a block size of 10k transactions. In the low-conflict P2P workload, Figure 10a, where $\alpha = 0.1$, iBTM scales from 12k-77k tps, reducing execution latency 3.20 \times . Block-STM achieves a comparable 73k tps at 32 threads. In high-conflict P2P, where $\alpha = 1.0$ (Figure 10a), throughput is modestly from 11k-14k tps. However, in Figure 10b, the batch transfer workload when $\alpha = 0.1$ exhibits better scalability, with iBTM increasing from 7k-50k tps, outperforming Block-STM's 47k tps. In the generic mixed workload with $\alpha = 0.1$ (Figure 10b) iBTM again scales well, improving from 9k-63k tps. Overall, iBTM consistently outperforms Block-STM and sequential execution, demonstrating the benefits of leveraging conflict specification for optimistic parallel execution.

7 Discussion and Concluding Remarks

The outstanding objective of this paper is to expose the significance of a execution methodology that prioritizes leveraging transactional conflict specifications as input for the parallel executor.

We demonstrated how this methodology can be constructive by implementing a state-of-the-art EVM parallel execution engine. We also demonstrated the methodology for MoveVM and its empirical benefits. We remark that we chose to exhibit this methodology for the read-write oblivious execution models of EVM and MoveVM when it would be more straightforward to implement our algorithms for the read-write aware execution model like Solana, since it does not require an explicit *conflict analyzer*. We leave that for a future iteration of the paper.

Localized Fee Markets. Our parallel execution methodology possibly lays the foundation for more localized dynamic *gas fee* marketplaces; if dependencies are specified a priori, transactions occurring in a congested contract of the blockchain state might be processed separately from others to prevent a localized state hotspot from increasing fees for the whole network. For example, a popular *non fungible token (NFT)* mint could create a large number of transaction requests in a short period of time [21]. A read-write aware blockchain can detect state hotspots upfront, such as the NFT minting event, to rate limit and charge a higher fee for transactions containing them [20, 34, 35]. This enables ordinary transactions to execute promptly, while transactions related to the minting process are prioritized independently based on the total gas associated with them and resulting congestion.

References

- 1 Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. Parblockchain: Leveraging transaction parallelism in permissioned blockchain systems. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 1337–1347, Los Alamitos, CA, USA, July 2019. IEEE, IEEE Computer Society. doi:10.1109/ICDCS.2019.00134.
- 2 Parwat Singh Anjana, Hagit Attiya, Sweta Kumari, Sathya Peri, and Archit Somani. Efficient concurrent execution of smart contracts in blockchains using object-based transactional memory. In *Networked Systems*, pages 77–93, Cham, 2021. Springer International Publishing. doi:10.1007/978-3-030-67087-0_6.
- 3 Parwat Singh Anjana, Sweta Kumari, Sathya Peri, Sachin Rathor, and Archit Somani. An efficient framework for optimistic concurrent execution of smart contracts. In *27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, pages 83–92. IEEE, IEEE Computer Society, February 2019. doi:10.1109/EMPDP.2019.8671637.
- 4 Parwat Singh Anjana, Sweta Kumari, Sathya Peri, Sachin Rathor, and Archit Somani. Optsmart: a space efficient optimistic concurrent execution of smart contracts. *Distributed and Parallel Databases*, pages 1–53, 2022. doi:10.1007/s10619-022-07412-y.
- 5 Aptos Labs. The aptos blockchain: Safe, scalable, and upgradeable web3 infrastructure, August 2022. [August 11, 2022]. URL: https://aptosfoundation.org/whitepaper/aptos-whitepaper_en.pdf.
- 6 Shrey Baheti, Parwat Singh Anjana, Sathya Peri, and Yogesh Simmhan. Dipetrans: A framework for distributed parallel execution of transactions of blocks in blockchains. *Concurrency and Computation: Practice and Experience*, 34(10):e6804, 2022. doi:10.1002/cpe.6804.
- 7 Block-stm implementation. [Github, Accessed: January 2024]. URL: <https://github.com/danielxiangz1/Block-STM>.
- 8 BlueAlloy. Revm: A rust implementation of the evm, 2023. [Github, Accessed: January 2025]. URL: <https://github.com/bluealloy/revm>.
- 9 Chainstack. Ethereum traceBlockByNumber API Reference, 2025. [Docs, Accessed: January 2025]. URL: <https://docs.chainstack.com/reference/ethereum-traceblockbynumber>.
- 10 CryptoKitties. CryptoKitties Webpage, 2024. [Webpage, Accessed: January 2025, Contract Address: 0x06012c8cf97BEaD5deAe237070F9587f8E7A266d, Creation Transaction: 0x691f348ef11e9ef95d540a2da2c5f38e36072619aa44db0827e1b8a276f120f4]. URL: <https://www.cryptokitties.co/>.

- 11 Thomas Dickerson, Paul Gazzillo, Maurice Herlihy, and Eric Koskinen. Adding concurrency to smart contracts. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 303–312, New York, NY, USA, 2017. PODC '17, ACM. doi:10.1145/3087801.3087835.
- 12 Ethereum (ETH): open-source blockchain-based distributed computing platform. [Webpage, Accessed: January 2025]. URL: <https://www.ethereum.org/>.
- 13 Ethereum Foundation. Ethereum 2.0 Merge, 2022. [Blog, Accessed: January 2025]. URL: <https://ethereum.org/en/upgrades/merge/>.
- 14 Etherscan verifier, July 2025. [Website, Accessed: July 2025]. URL: <https://etherscan.io/verifyContract>.
- 15 Mohamed Fouda. The case for parallel processing chains, September 2022. [Blog, Accessed: January 2025]. URL: <https://medium.com/alliancedao/the-case-for-parallel-processing-chains-90bac38a6ba4>.
- 16 Sui Foundation. All about parallelization, January 2024. [Blog, Accessed: January 2025]. URL: <https://blog.sui.io/parallelization-explained/>.
- 17 Rati Gelashvili, Alexander Spiegelman, Zhuolun Xiang, George Danezis, Zekun Li, Dahlia Malkhi, Yu Xia, and Runtian Zhou. Block-stm: Scaling blockchain execution by turning ordering curse to a performance blessing. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, PPOPP '23, pages 232–244, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3572848.3577524.
- 18 Rachid Guerraoui and Michal Kapalka. *Principles of Transactional Memory, Synthesis Lectures on Distributed Computing Theory*. Morgan and Claypool, 2010.
- 19 Petr Kuznetsov and Srivatsan Ravi. On the cost of concurrency in transactional memory. In *International Conference on Principles of Distributed Systems (OPODIS)*, pages 112–127, 2011. doi:10.1007/978-3-642-25873-2_9.
- 20 Helius Labs. Solana local fee markets: How they work and why they matter, 2024. [Blog, Accessed: May 2025]. URL: <https://www.helius.dev/blog/solana-local-fee-markets>.
- 21 What is lazy minting? introducing a smarter yet economical way to mint nfts, January 2024. [Blog, Accessed: January 2025]. URL: <https://www.antiersolutions.com/what-is-lazy-minting-introducing-a-smarter-yet-economical-way-to-mint-nfts/>.
- 22 Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- 23 Mythx vulnerability coverage, July 2025. [Website, Accessed: July 2025]. URL: <https://mythx.io/detectors/>.
- 24 Pareto distribution, 2023. [Wikipedia, Accessed: January 2025]. URL: https://en.wikipedia.org/wiki/Pareto_distribution.
- 25 Manaswini Piduguralla, Saheli Chakraborty, Parwat Singh Anjana, and Sathya Peri. Dag-based efficient parallel scheduler for blockchains: Hyperledger sawtooth as a case study. In *European Conference on Parallel Processing*, pages 184–198, Berlin, Heidelberg, 2023. Springer, Springer-Verlag. doi:10.1007/978-3-031-39698-4_13.
- 26 Polygon Labs. Innovating the main chain: A polygon pos study in parallelization, 2024. [Blog, Accessed: May 2025]. URL: <https://polygon.technology/blog/innovating-the-main-chain-a-polygon-pos-study-in-parallelization>.
- 27 RISE Chain. RISE Chain Webpage, 2024. [Webpage, Accessed: January 2025]. URL: <https://www.riselabs.xyz/>.
- 28 RISE Labs. PEVM: Parallel Ethereum Virtual Machine, 2023. [Github, Accessed: January 2025]. URL: <https://github.com/risechain/pevm>.
- 29 Vikram Saraph and Maurice Herlihy. An empirical study of speculative concurrency in ethereum smart contracts. In *International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2019)*, pages 4:1–4:15, Dagstuhl, Germany, 2019. OpenAccess Series in Informatics (OASICS), Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/OASICS.Tokenomics.2019.4.

- 30 Architecture documentaion v1.2. [Whitepaper, Accessed: January 2025]. URL: https://sawtooth.hyperledger.org/docs/1.2/architecture/transaction_scheduling.html.
- 31 Sei Labs. Sei Giga: Scaling the EVM through Parallel Execution, 2024. [Whitepaper, Accessed: May 2025]. URL: https://github.com/sei-protocol/sei-chain/blob/main/whitepaper/Sei_Giga.pdf.
- 32 Nir Shavit and Dan Touitou. Software transactional memory. In *Proceedings of the fourteenth annual ACM symposium on Principles of distributed computing*, pages 204–213, 1995. doi: 10.1145/224964.224987.
- 33 Solana documentation. [Docs, Accessed: January 2025]. URL: <https://docs.solana.com/>.
- 34 Solana fees, part 1, December 2023. [Blog, Accessed: May 2025]. URL: <https://www.umbraresearch.xyz/writings/solana-fees-part-1>.
- 35 SolanaFloor. Solana’s local fee market: A solution to soaring gas prices, 2024. [Blog, Accessed: May 2025]. URL: <https://solanafloor.com/news/solanas-local-fee-market-a-solution-to-soaring-gas-prices>.
- 36 Sui documentation: Discover the power of sui through examples, guides, and concepts. [Docs, Accessed: January 2025]. URL: <https://docs.sui.io>.
- 37 Umbraresearch. Lifecycle of a solana transaction. [Blog, Accessed: January 2025]. URL: <https://www.umbraresearch.xyz/writings/lifecycle-of-a-solana-transaction>.
- 38 Anatoly Yakovenko. Sealevel - parallel processing thousands of smart contracts, September 2019. [Blog, Accessed: January 2025]. URL: <https://medium.com/solana-labs/sealevel-parallel-processing-thousands-of-smart-contracts-d814b378192>.