Efficient Dynamic Concurrency Analysis with Collective Sparse Segment Trees

HÜNKAR CAN TUNÇ, Computer Science, Aarhus Universitet, Aarhus, Denmark
YIFAN DONG, Computer Science, Aarhus Universitet, Aarhus, Denmark
AMEYA PRASHANT DESHMUKH, IIT Bombay, Mumbai, India
BERK CIRISCI, Amazon Web Services Germany GmbH, Berlin, Germany
CONSTANTIN ENEA, Université de Paris, Paris, France
ANDREAS PAVLOGIANNIS, Computer Science, Aarhus Universitet, Aarhus, Denmark

Dynamic analyses are a standard approach to analyzing and testing concurrent programs. Such techniques observe program traces σ and analyze them to infer the presence or absence of bugs. At its core, each analysis maintains a partial order P that represents order dependencies between the events of σ. Naturally, the scalability of the analysis largely depends on maintaining P efficiently. The standard data structure for this task has thus far been Vector Clocks. These, however, are slow for analyses that follow a non-streaming style, costing $O(n)$ time for inserting (and propagating) each new ordering in P, where n is the size of σ, while they cannot handle the deletion of existing orderings.

In this article, we develop *Collective Sparse Segment Trees* (*CSSTs*), a simple but elegant data structure for maintaining a partial order P. CSSTs thrive when the width k of P is much smaller than the size n of its domain, allowing inserting, deleting, and querying for orderings in P to run in $O(\log n)$ time. For a concurrent trace, k normally equals the number of its threads, and is orders of magnitude smaller than its size n, making CSSTs fitting for this setting. Our experiments confirm that CSSTs are the best data structure currently to handle a range of dynamic analyses from existing literature.

CCS Concepts: • **Software and its engineering → Software verification and validation**; • **Theory of computation** → *Theory and algorithms for application domains*; *Program analysis;*

Additional Key Words and Phrases: Concurrency, happens-before, vector clocks, dynamic concurrency analyses, dynamic reachability

This article, originally titled "CSSTs: A Dynamic Data Structure for Partial Orders in Concurrent Execution Analysis", was invited by ACM TOCS for extended publication through recommendation by the Chairs of ASPLOS 2024.
This work was partially supported by a research grant (VIL42117) from VILLUM FONDEN, and by a research grant from STIBOFONDEN..
Authors' Contact Information: Hünkar Can Tunç (corresponding author), Computer Science, Aarhus Universitet, Aarhus, Denmark; e-mail: hcantunc@gmail.com; Yifan Dong, Computer Science, Aarhus Universitet, Aarhus, Denmark; e-mail: yvan.dong@cs.au.dk; Ameya Prashant Deshmukh, IIT Bombay, Mumbai, India; e-mail: meyapd@cse.iitb.ac.in; Berk Cirisci, Amazon Web Services Germany GmbH, Berlin, Germany; e-mail: cirisci@amazon.de; Constantin Enea, Université de Paris, Paris, France; e-mail: cenea@irif.fr; Andreas Pavlogiannis, Computer Science, Aarhus Universitet, Aarhus, Denmark; e-mail: pavlogiannis@cs.au.dk.

ACM 0734-2071/2026/08-ART25
https://doi.org/10.1145/3773085

ACM Reference Format:
Hünkar Can Tunç, Yifan Dong, Ameya Prashant Deshmukh, Berk Cirisci, Constantin Enea, and Andreas Pavlogiannis. 2026. Efficient Dynamic Concurrency Analysis with Collective Sparse Segment Trees. *ACM Trans. Comput. Syst.* 44, 3, Article 25 (August 2026), 31 pages. https://doi.org/10.1145/3773085

1 Introduction

Concurrent programming is notoriously difficult due to the inherent nondeterminism in interprocess communication [30]. As such, considerable efforts go toward effective techniques for analyzing concurrent programs and testing them for correctness. Dynamic analyses are one of the most widespread types of such techniques. Their main working principle is to discover software faults by analyzing concrete program executions. Prominent examples include concurrency bugs such as data races [14, 16, 18, 24, 25, 27, 32–34], deadlocks [9, 21, 41], linearizability [13] and atomicity violations [7, 9, 17, 25], thread-safety [23], and memory violations [20, 42], to name a few.

Although each analysis employs reasoning that is specific to the problem at hand, one of its core components is nearly always the construction of a partial order P (often termed generically as "happens-before" [22]) that captures order dependencies between events of the analyzed trace σ. The analysis undergoes a sequence of interleaved operations of (i) *inserting* new orderings $e_1 \to e_2$ in P, (ii) *querying* whether $e_1 \to e_2$ in P, and even (iii) *deleting* existing orderings $e_1 \to e_2$ from P, as it attempts to prove that σ has certain properties (e.g., that it is sequentially-consistent, or it contains a data race). To benefit scalability, each operation must take as little time as possible, as both the number of traces and the size of each trace normally span several orders of magnitude.

Consider analyzing a trace σ of size n. Representing P as a graph G, the above setting is colloquially known as *dynamic graph reachability*, with dynamically inserting/deleting edges and performing reachability queries. Such queries are handled in $O(1)$ time, at the cost of $O(n^2)$ for keeping G transitively closed after each update. This bound is prohibitively large for typical values of n in the range 10^4-10^{10}.

The scalability of dynamic analyses has been driven primarily by a simple data structure known as *Vector Clocks* [29], based on the realization that (i) σ consists of a small number of threads k, and (ii) normally, P contains k totally ordered chains (or a number proportional to k). Consequently, the whole backward set of an event e can be compactly summarized as an integer array, i.e., a Vector Clock. This allows to replace one factor n in the update complexity by the much smaller k, i.e., a single edge insertion now takes $O(nk)$ time.

In many analyses that have a *streaming nature*, the insertion cost of Vector Clocks is further reduced to $O(k)$. Here the events of σ are processed one by one, and new orderings $e_1 \to e_2$ are always inserted where e_2 is the current event being processed. Popular examples include analyses for detecting data-races [16, 25, 34], deadlocks [41], and atomicity violations [28], as well as consistency checking in some weak-memory models [39]. Thus in the streaming setting, Vector Clocks are arguably the best data structure to represent a partial order. On the other hand, many analyses are *non-streaming*, i.e., as the analysis progresses, it might insert new orderings between arbitrary events of σ. These new orderings might have to be propagated across n events, recovering the $O(nk)$ bound, which can lead to a significant slowdown. There are numerous analyses in practice that operate in this non-streaming manner. Examples include consistency checking under various memory models [8, 35], predictive analyses of concurrency bugs [9, 20, 32, 42], and testing database isolation levels [6]. A natural question thus arises: *can the $O(nk)$ cost/operation for non-streaming analyses be reduced further?* In general, *what is the best data structure for non-streaming concurrent execution analysis?*

Here we develop a simple but elegant data structure for this general setting, called ***Collective Sparse Segment Trees*** (CSSTs). Like Vector Clocks, CSSTs exploit the fact that $k \ll n$, and offer

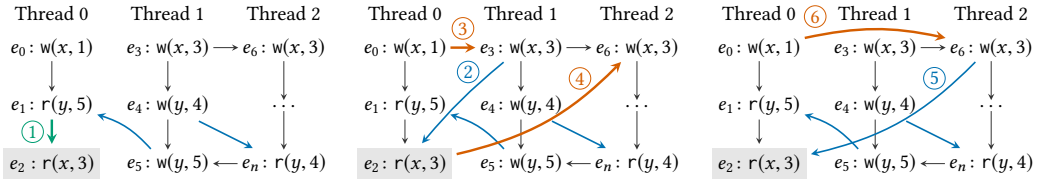

(a) Initial partial order. (b) First update. (c) Delete and update.

Fig. 1. A consistency analysis. Blue edges represent reads-from information. Numbered edges represent new orderings inserted by the analysis.

Table 1. Update and Query Times for Maintaining a Partial Order of n Events, k Chains, and Maximum Degree δ

	Insert	**Delete**	**Query**
`Graphs`	$O(1)$	$O(1)$	$O(n)$
`VCs`	$O(nk)$	Not supported	$O(1)$
`CSSTs`	$O(\max(\log \delta, \min(\log n, d)))$	$O(\max(\log \delta, \min(\log n, d)))$	$O(k^3 \min(\log n, d))$
`CSSTs` (incremental)	$O(k^2 \min(\log n, d))$	Not supported	$O(\min(\log n, d))$

The cross-chain density parameter d is introduced here and captures a type of sparsity prevalent in practice. Note that all times for CSSTs are bounded by $O(\log n)$.

update and querying operations that run roughly in $O(\log n)$ time, thereby reducing the prohibitive linear dependency on n to logarithmic. We argue that, as Vector Clocks are the most fitting data structure for handling *streaming* dynamic analyses, `CSSTs` are the most fitting data structure for handling non-streaming analyses. Table 1 shows a summary of the running times achieved by each standard data structure (`Graphs` and `VCs`), as well as by `CSSTs` introduced in this work (both in their fully dynamic and only incremental setting).

1.1 Motivating Example

Consider the analysis of an abstract trace shown in Figure 1, consisting of three threads and write/read events $\mathtt{w}(x, v)/\mathtt{r}(x, v)$, where x is a variable and v is a value. The consistency-testing problem asks whether there is an interleaving that is consistent with the values observed by each read [19], and has numerous applications in dynamic concurrency analyses [4, 8, 10, 11, 15, 20, 21, 24, 35, 41]. The analysis maintains (i) a reads-from map $\mathtt{w}(x, v) \mapsto \mathtt{r}(x, v)$, denoting that $\mathtt{r}(x, v)$ obtains its value from $\mathtt{w}(x, v)$, and (ii) a partial order P that is induced by this reads-from map and any additional indirect orderings this might impose.

Assume that the consistency analysis has, so far, established the partial order P in Figure 1(a), with blue edges marking the current reads-from map. The analysis proceeds on the read event $e_2 : \mathtt{r}(x, 3)$. It first inserts in P the program order edge (1). As e_2 has no successors in P, this edge need not be propagated further. Now the analysis must decide which of the two writes e_3 and e_6 can be observed by e_2. In general, there is no efficient procedure for deciding this [19], and the consistency algorithm proceeds by trying each option.

In Figure 1(b), the analysis tries to map $e_3 \mapsto e_2$, leading to inserting (2) in P. Moreover, since $e_0 \to e_2$ in P, the analysis also inserts (3) and (4) in P. These are necessary orderings to respect the fact that e_2 reads from e_3 — the process of inferring such orderings is known as saturation, and is used widely in dynamic analyses (e.g., [2, 24, 32, 34, 35, 43, 44]). Crucially, both e_3 and e_6 have many

successors in P, thus these edges must be propagated along many events, costing $O(n)$ time, even if P is represented using Vector Clocks. In contrast, CSSTs achieve a much smaller $O(\log n)$ bound.

The reads-from choice $e_3 \mapsto e_2$ is inconsistent due to the cycle $e_2 \to e_6 \to \ldots \to e_n \to e_5 \to e_1 \to e_2$. The analysis must delete the orderings (2), (3) and (4), before proceeding with an alternative writer for e_2. When P is represented via Vector Clocks, there is no efficient method for deletion, and essentially it must be recomputed from scratch, taking $O(n^2)$ time. In contrast, CSSTs only need $O(\log n)$ time per edge deletion, thereby completing this task efficiently.

Finally, in Figure 1(c), the analysis tries to map $e_6 \mapsto e_2$, leading to inserting (5) in P. It also infers ordering (6), which again must be propagated further, a task that costs $O(n)$ time for Vector Clocks but only $O(\log n)$ time for CSSTs.

1.2 Contributions

We develop CSSTs, a new data structure for maintaining and querying partial orders of small width. In the concurrent execution analysis setting, the width of the partial order is typically bounded by the number of threads in the execution. Concretely, consider the maintenance of a DAG G, representing a partial order P over n events and k totally ordered chains (normally k equals the number of threads).

1. CSSTs support fully dynamic updates (inserting and deleting edges) in $O(\max(\log \delta, \min(\log n, d)))$ time, where δ is the maximum degree in G and d is the *cross-chain density* of G, a novel notion we introduce here that captures a type of sparsity prevalent in practice. Intuitively, the cross-chain density corresponds to the maximum, over all chains, of the number of nodes within a chain that have outgoing edges to nodes in other chains (considering only explicit edges and excluding transitive orderings). Normally, δ is constant and the above expression results in $O(d)$. Reachability (ordering) queries take $O(k^3 \min(\log n, d))$ time.
2. We further optimize CSSTs for the purely incremental setting, where edges can be inserted but not removed. Incremental CSSTs run updates and queries in time $O(k^2 \min(\log n, d))$ and $O(\min(\log n, d))$, respectively,
3. We make an extensive experimental evaluation of CSSTs in concurrency analyses from the literature, spanning data-race detection [32], deadlock detection [9], memory bugs [42], consistency checking for X86-TSO [35], use-after-free bugs [20], data-races in the C11 memory model [24], view serializability [6], and root cause analysis of linearizability violations [13]. Our results show that CSSTs speed up the corresponding analysis in most cases, indicating that they are the most fitting data structure in this ubiquitous setting.

CSSTs are based on the insight that dynamic reachability on partial orders of small width can be reduced to a basic algorithmic problem known as *dynamic suffix minima*. This idea was partly explored in a data structure for incremental reachability underpinning the M2 data-race detector [32]. However, CSSTs are more advanced, as (i) they employ techniques making them faster in all cases, (ii) they achieve tighter time and space bounds on sparse instances, and (iii) they support fully dynamic updates (i.e., both incremental and decremental). Our experimental results confirm that CSSTs are indeed more efficient in both time and memory, besides being broader (as they handle edge deletions).

CSSTs are available in a repository [1] as a stand alone data structure for dynamic graph reachability. An earlier version of this work appeared in [40].

2 Preliminaries

In this section, we develop some general notation and introduce standard concepts on concurrent executions and partial orders that will be used throughout the article.

2.1 Concurrent Execution Model

Events and traces. We broadly consider traces σ of concurrent programs, representing the history of operations in an execution. We write Events_σ to denote the set of events occurring in σ. Each event of σ is a tuple $e = \langle t, i, m\rangle$, where t is an identifier of the thread that performed e and i is the sequence id of e. The pair $\langle t, i\rangle$ serves as a unique identifier for e. The field m represents meta information for e that might be relevant to a dynamic analysis, but is immaterial for CSSTs. E.g., m may record the operation performed by e (read, write, etc), as well as the variable it accesses and the value it acquires. CSSTs operate only on the identifier $\langle t, i\rangle$, while the meta information m will be used only in our examples. We write $\text{tid}(e)$, $\text{id}(e)$ for the thread identifier and sequence id of e, respectively. Depending on the application, the events in σ may or may not be totally ordered.

Relations and functions on execution traces. For a binary relation X, we denote by X^* its reflexive transitive closure. A partial order is a reflexive, transitive, and anti-symmetric binary relation defined on the elements of a given set S. We denote by $\leq^\sigma_{\mathsf{P}}$ a partial order P over the set Events_σ. Given $e_i, e_j \in \text{Events}_\sigma$, we write $e_i \leq^\sigma_{\mathsf{P}} e_j$ to represent $(e_i, e_j) \in \leq^\sigma_{\mathsf{P}}$. We write $e_i <^\sigma_{\mathsf{P}} e_j$ to represent $e_i \leq^\sigma_{\mathsf{P}} e_j$ and $e_i \neq e_j$. The *program order* $\leq^\sigma_{\text{po}}$ represents a partial order where $e_i \leq^\sigma_{\text{po}} e_j$ iff $\text{id}(e_i) \leq \text{id}(e_j)$ and $\text{tid}(e_i) = \text{tid}(e_j)$. We write $\texttt{width}(P)$ to represent the largest size of a set $S \subseteq \text{Events}_\sigma$ such that for all distinct pairs $e_i, e_j \in S$, we have that $e_i \not\leq^\sigma_{\mathsf{P}} e_j$ and $e_j \not\leq^\sigma_{\mathsf{P}} e_i$.

In concurrency analyses, $\texttt{width}(P)$ is almost always bounded by some constant, a fact that is exploited by Vector Clocks, as well as the CSSTs we introduce in this work. E.g., in many analyses $\leq^\sigma_{\text{po}} \subseteq \leq^\sigma_{\mathsf{P}}$, implying that $\texttt{width}(P) \leq k$, where k is the number of threads. Although in other settings (e.g., those involving weak-memory concurrency) $\texttt{width}(P)$ might be larger than k, it still remains bounded by some small factor of k. For example, analyses involving TSO normally satisfy $\texttt{width}(P) \leq 2k$ (we touch on this later in Section 5).

2.2 Dynamic Reachability on DAGs

Chain directed acyclic graphs. We consider **directed acyclic graphs** (**DAGs**) $G = (V, E)$, where V is the set of nodes and E is the set of edges. Dynamic analyses use such DAGs to represent partial orders on n events of a concurrent execution, where k is normally the number of threads, or a value proportional to that. For this purpose, a node in V is a pair $u = \langle t, i\rangle \in [k] \times [n]$, where $[j] = \{0, 1, \ldots, j-1\}$. The set of edges normally captures the program order, which is reflected in the fact that for any two nodes $\langle t, i\rangle, \langle t, i+1\rangle \in V$, we have $(\langle t, i\rangle, \langle t, i+1\rangle) \in E$. We can thus represent G as k totally-ordered chains with additional edges across chains, and call such graphs *chain DAGs*. A node $\langle t, i\rangle$ is said to be *earlier* than another node in the same chain $\langle t, j\rangle$ if $i \leq j$, and *later* otherwise. Then, G can be written as the composition of $k(k-1)$ subgraphs $\texttt{Sub(G)}^{t_2}_{t_1}$ induced by the chains t_1 and t_2, as well as the cross-chain edges from t_1 to t_2 (see Figure 2). The *cross-chain density* d of G is the maximum, among all chains t_1, number of nodes in chain t_1 that have an outgoing edge to some other chain.

Dynamic reachability. The dynamic reachability problem is defined with respect to a given chain DAG $G = (V, E)$ and online sequence of operations of the following types.

1. An $\texttt{insertEdge}(u, v)$ operation, such that $(u, v) \notin E$, inserts the edge (u, v) in G.
2. A $\texttt{deleteEdge}(u, v)$ operation, such that $(u, v) \in E$, deletes the edge (u, v) in G.
3. A $\texttt{reachable}(u, v)$ operation returns True iff $u \rightarrow^* v$.
4. A $\texttt{successor}(u, t)$ operation returns the earliest successor of u in the tth chain.
5. A $\texttt{predecessor}(u, t)$ operation returns the latest predecessor of u in the tth chain.

The operations insertEdge and deleteEdge are updates, while reachable, successor , and predecessor are queries. We allow update operations only across nodes in different chains. In this

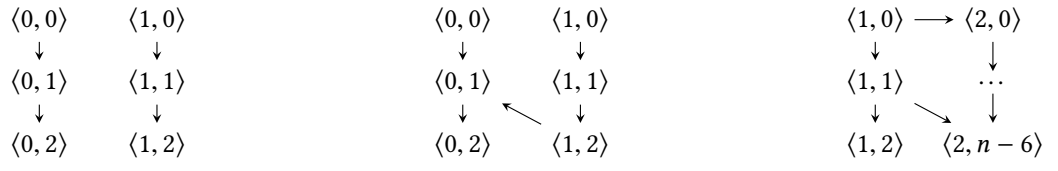

(a) $\mathsf{Sub}(\mathsf{G})^1_0$. (b) $\mathsf{Sub}(\mathsf{G})^0_1$. (c) $\mathsf{Sub}(\mathsf{G})^2_1$.

Fig. 2. Three subgraphs of the chain DAG G formed from the execution in Figure 1(a). The cross-chain density of G is 2, as chain 1 (i.e., thread 1 in Figure 1(a)) has two outgoing edges to chain 2 (from events e_3 and e_4).

dynamic setting, we define the cross-chain density of G to be the maximum cross-chain density of the graph across all update operations on G.

3 Collective Sparse Segment Trees

We now present CSSTs, a data structure for efficiently solving the dynamic reachability problem in the context of concurrent execution analysis. The properties of CSSTs are stated in the following theorem.

Theorem 1. *Consider a chain DAG G of n nodes, k chains, maximum out-degree δ and cross-chain density d.* CSSTs *maintain G under dynamic updates, costing $O(\max(\log\delta, \min(\log n, d)))$ time per update and $O(k^3 \min(\log n, d))$ time per query.*

CSSTs are based on the insight that dynamic reachability on DAGs with few chains can be efficiently reduced to repeated instances of another basic problem known as *dynamic suffix minima*. That is, by exploiting the chain decomposition of the DAG, we map dynamic reachability problem to dynamic suffix minima problem within each pair of chains. We begin by developing this concept when G consists of only $k = 2$ chains in Section 3.1. Dynamic suffix minima are solved using a data structure known as ***Segment Trees (STs)***. A key novelty of CSSTs is based on the realization that most direct orderings in dynamic analyses make G have low cross-chain density (i.e., the cross-chain edges are sparse). We exploit this insight in Section 3.2, by developing ***Sparse Segment Trees*** (**SSTs**) that handle arbitrary graphs, but become more efficient the lower their cross-chain density becomes. Finally, in Section 3.3 we handle DAGs of arbitrarily many chains k, by appropriately maintaining *collections* of SSTs, which also explains the name of the data structure.

3.1 Dynamic Suffix Minima and Segment Trees

Here we state the algorithmic problem of dynamic suffix minima, and relate it to dynamic reachability on chain DAGs, focusing on the special case of $k = 2$ chains. In later sections we explain how to handle $k > 2$. Some insights are based on observations made in [32] for the purely incremental case.

Dynamic suffix minima. The dynamic suffix minima problem concerns maintaining an array $\mathtt{A}$ of n values in $\mathbb{N} \cup \{\infty\}$ under an online sequence of updates, and answering minima queries on ranges of $\mathtt{A}$. We write $\mathtt{A}[i]$ and $\mathtt{A}[i : j]$ to denote the value of $\mathtt{A}$ at index i and the sequence of values $\mathtt{A}[i], \ldots, \mathtt{A}[j]$, respectively, and write $\mathtt{A}[i :]$ as shorthand for $\mathtt{A}[i : |\mathtt{A}| - 1]$. We say that $\mathtt{A}$ is empty on index i if $\mathtt{A}[i] = \infty$. The density of $\mathtt{A}$ is the number of non empty entries of $\mathtt{A}$.

For $i \in [n]$ and $a \in \mathbb{N} \cup \{\infty\}$, an operation on $\mathtt{A}$ is one of:

1. $\mathtt{min}(\mathtt{A}, i)$, returning $\mathtt{min}_{i \le j < n} \mathtt{A}[j]$, i.e., the minimum value in the suffix $\mathtt{A}[i :]$;
2. $\mathtt{argleq}(\mathtt{A}, a)$, returning $\max_{i\,:\,\mathtt{A}[i] \le a} i$, i.e., the largest index i of $\mathtt{A}$ such that $\mathtt{A}[i] \le a$; and
3. $\mathtt{update}(\mathtt{A}, i, a)$, setting $\mathtt{A}[i] = a$.

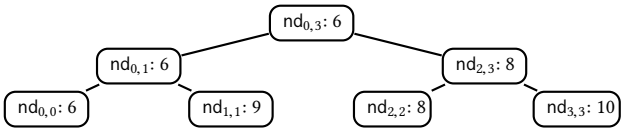

Fig. 3. A segment tree representation of $A = [6, 9, 8, 10]$.

The task is to answer `min` and `argleq` queries, taking into consideration all preceding `update` operations.

Dynamic suffix minima (in fact, the more general problem of dynamic *range* minima) is solved efficiently using STs, taking $O(\log n)$ time per operation (see, e.g., [5]). A ST T represents A as a binary tree of nodes $\mathsf{nd}_{i,j}$ each associated with a range $[i, j]$. Each $\mathsf{nd}_{i,j}$ maintains a value $\mathsf{nd}_{i,j}.\mathtt{min} = \min(A[i : j])$. We write $\mathtt{depth}(\mathsf{nd}_{i,j})$ to denote the depth of $\mathsf{nd}_{i,j}$ in T. We let $T.\mathtt{root}$ be the root node of T. Each query is answered by traversing T from the root to a leaf, the latter containing the value to be returned.

Example 1. Consider the ST T in Figure 3 representing the array $A = [6, 9, 8, 10]$. We have the following.
— $\mathtt{min}(A, 0) = 6$, $\mathtt{min}(A, 1) = \mathtt{min}(A, 2) = 8$, $\mathtt{min}(A, 3) = 10$
— $T.\mathtt{root}.\mathtt{min} = 6$, $\mathsf{nd}_{2,3}.\,\mathrm{min} = 8$
— $\mathtt{argleq}(A, 7) = 0$, $\mathtt{argleq}(A, 9) = 2$, $\mathtt{argleq}(A, 11) = 3$
— $\mathtt{update}(A, 3, 7)$ sets $A[3] = 7$. In T, this is captured by setting $\mathsf{nd}_{2,3}.\mathtt{min} = \mathsf{nd}_{3,3}.\mathtt{min} = 7$.

Dynamic reachability on $k = 2$ chains. Recall that a chain DAG G of k chains is the composition of $k(k-1)$ subgraphs $\mathsf{Sub}(\mathsf{G})_{t_1}^{t_2}$ (Figure 2). Now assume that $k = 2$. We represent each $\mathsf{Sub}(\mathsf{G})_{t_1}^{t_2}$ as an array $\mathsf{A}_{t_1}^{t_2} : [n] \to [n] \cup \infty$. $\mathsf{A}_{t_1}^{t_2}[j_1]$ stores a unique neighbor of node $\langle t_1, j_1 \rangle$ in chain t_2. What if $\langle t_1, j_1 \rangle$ has multiple neighbors to t_2? Then, it suffices to store the *earliest neighbor* of $\langle t_1, j_1 \rangle$ in t_2. Although this is not an exact representation of G, it captures exactly all reachability in G. In particular, we maintain the invariant

$$\mathsf{A}_{t_1}^{t_2} = \min(\{j_2 : \langle t_1, j_1 \rangle \to \langle t_2, j_2 \rangle\}) \tag{1}$$

under the standard convention that $\min(\emptyset) = \infty$. A query $\mathtt{successor}(\langle t_1, j_1 \rangle, t_2)$ simply returns the suffix minima $x = \mathtt{min}(\mathsf{A}_{t_1}^{t_2}, j_2)$, while a query $\mathtt{predecessor}(\langle t_2, j_2 \rangle, t_1)$ returns the suffix arg-minima $\mathtt{argleq}(\mathsf{A}_{t_1}^{t_2}, j_2)$. For reachability queries, observe that we have $\langle t_1, j_1 \rangle \to^* \langle t_2, j_2 \rangle$ in G iff there is an edge $\langle t_1, i_1 \rangle \to \langle t_2, i_2 \rangle$ such that $i_1 \geq j_1$ and $i_2 \leq j_2$. Hence, a query $\mathtt{reachable}(\langle t_1, j_1 \rangle, \langle t_2, j_2 \rangle)$ reduces to checking whether $\mathtt{successor}(\langle t_1, j_1 \rangle, t_2) \leq j_2$.

How can we maintain the invariant Equation (1) under edge insertions and deletions? We achieve this by storing all current successors of $\langle t_1, j_1 \rangle$ to chain t_2 in a min heap $\mathtt{edgeHeap}_{t_1}^{t_2}[j_1]$. Inserting and deleting edges from $\langle t_1, j_1 \rangle$ is handled by applying the same operation in $\mathtt{edgeHeap}_{t_1}^{t_2}[j_1]$. This ensures that the earliest neighbor of $\langle t_1, j_1 \rangle$ in t_2 appears in the root of $\mathtt{edgeHeap}_{t_1}^{t_2}[j_1]$, which is also stored in $\mathsf{A}_{t_1}^{t_2}[j_1]$.

Example 2. Consider again the partial order $\leq_{\mathrm{P}}^{\sigma}$ in *Figure* **??**. Some of the arrays $\mathsf{A}_{t_1}^{t_2}$ are shown in Figure 4. Observe that $\mathtt{successor}(e_3, 2) = e_6$ can be determined by the suffix minima query $\mathtt{min}(\mathsf{A}_1^2, 0)$. Similarly, $\mathtt{predecessor}(e_2, 1) = e_5$ can be determined by the suffix arg-minima query $\mathtt{argleq}(\mathsf{A}_1^0, 1)$.

(a) Array A_0^1. (b) Array A_1^0. (c) Array A_1^2. (d) Array A_2^1.

Fig. 4. Representation of the partial order in Figure 1(a).

Handing $k > 2$ chains. Naturally, when $k > 2$, we can have a suffix minima array $\mathsf{A}_{t_1}^{t_2}$ for every pair of threads $t_1, t_2 \in [k]$ with $t_1 \neq t_2$. However, a path $\langle t_1, j_1 \rangle \rightarrow^* \langle t_2, j_2 \rangle$ might go through intermediate nodes $\langle t_3, j_1 \rangle$, i.e., reachability is formed transitively across chains (e.g., the path $e_6 \rightarrow^* e_2$ in Figure 1(a)). We address this challenge depending on the setting (fully dynamic or purely incremental).

On a high level, in the fully dynamic setting (Section 3.3) in each $\mathsf{A}_{t_1}^{t_2}$ we store direct edges, i.e., each edge insertion is handled similarly to what was already described for $k = 2$. During queries, we perform a transitive closure operation across chains. Taking advantage of the properties of our data structure, this transitive closure takes $O(k^3)$ time, as opposed to $O(n^2)$ time, and thus remains efficient.

In the purely incremental setting (Section 4), on the other hand, $\mathsf{A}_{t_1}^{t_2}$ stores transitive reachability across chains, which is computed every time a new edge inserted. Again, taking advantage of the properties of our data structure, this transitive closure takes $O(k^2)$ time, as opposed to $O(n)$ time, and is thus very efficient. Now, queries can be performed by directly querying the respective suffix minima array, similarly to what was already described for $k = 2$.

3.2 Sparse Segment Trees

Using STs T to solve suffix minima, each operation completes in time proportional to the height of T, i.e., $O(\log n)$. In practice n can be very large and thus hinder performance. Here we develop SSTs, that optimize STs for our setting. These are based on two main ideas, namely (i) *minima indexing*, and (ii) *sparse tree representation*.

Intuitively, STs solve the more general problem of *range minima queries*, where every query asks for the minima in a range $A[i : j]$, as opposed to the suffix $A[i :]$. Minima indexing exploits the fact that our queries always concern a suffix of A, and allows them to often complete without traversing a full path from the root to a leaf.

Sparse tree representation is an optimization of the tree representation, which becomes more condensed the lower the density of A is (i.e., the fewer non-∞ entries it has), stemming from the observation that ∞ entries do not contribute to the queries. In a dynamic analysis setting, G normally has low cross-chain density d, which means that the suffix-minima arrays storing reachability across chains (as described in Section 3.1) are sparse. This reduces the height of the tree, allowing both queries and updates to run faster.

Minima indexing. Each node $\mathsf{nd}_{i,j}$ of a SST T maintains a field $\mathtt{pos} \in [i, j]$, that is defined recursively: it points to the largest index of A that contains the minimum element in the range $A[i : j]$, after excluding all indexes that are stored in the pos fields of all ancestors of $\mathsf{nd}_{i,j}$. Formally, let $\mathsf{nd}^1, \ldots, \mathsf{nd}^m$ be the ancestors of $\mathsf{nd}_{i,j}$ in T. Then

$$\mathsf{nd}_{i,j}.\mathtt{pos} = \max \left(\operatorname*{arg\,min}_{\ell \in [i,j] \setminus \{\mathsf{nd}^1.\mathtt{pos}, \ldots, \mathsf{nd}^m.\mathtt{pos}\}} A[\ell] \right) \tag{2}$$

We also set $\mathsf{nd}_{i,j}.\mathtt{min} = A[\mathsf{nd}_{i,j}.\mathtt{pos}]$. Observe that now $\mathsf{nd}_{i,j}.\mathtt{min}$ might not be the minimum value in $A[i : j]$, in contrast to the standard STs.

(a) Array A.

(b) ST representation of $\mathtt{A}_{t_1}^{t_2}$.

Fig. 5. ST representation with minima indexing. The entries $A[4:6]$ are > 59. The notation $\mathsf{nd}_{i,j} : (x, y)$ denotes $\mathsf{nd}_{i,j}.\mathtt{min} = x$, $\mathsf{nd}_{i,j}.\mathsf{pos} = y$.

The intuition behind $\mathsf{nd}_{i,j}.\mathsf{pos}$ is as follows. Each query $\mathtt{min}(A, i)$ starts a traversal from $T.\mathtt{root}$ to a leaf. If we encounter a node $\mathsf{nd}_{i,j}$ with $\mathsf{nd}_{i,j}.\mathsf{pos} \geq i$, the traversal can stop early and return that $\mathtt{min}(A, i) = \mathsf{nd}_{i,j}.\mathtt{min}$. Since we always query on a suffix of A, storing the *largest* index pointing to a minima increases the chances of stopping early. Moreover, we exclude the pos fields of the ancestors of $\mathsf{nd}_{i,j}$ as these are irrelevant for $\mathsf{nd}_{i,j}$, in the sense that this index would have already stopped the traversal in an ancestor of $\mathsf{nd}_{i,j}$.

Example 3. Consider the array A in Figure 5(a), and part of the corresponding ST in Figure 5(b). We have $\mathsf{nd}_{0,7}.\mathsf{pos} = 1 = \arg\min_i A[i]$. A query $\mathtt{min}(A, 0)$ is now directly answered at the root, as $\mathsf{nd}_{0,7}.\mathsf{pos} \in [0:7]$. On the other hand, the query $\mathtt{min}(A, 2)$ is not answered at the root as $\mathsf{nd}_{0,7}.\mathsf{pos} \notin [2:7]$. The procedure then traverses its child $\mathsf{nd}_{0,3}$. We have $\mathsf{nd}_{0,3}.\mathsf{pos} = 3$, which indexes the smallest element in $A[0:3]$ after excluding $\mathsf{nd}_{0,7}.\mathsf{pos}$. The traversal stops here, as $\mathsf{nd}_{0,3}.\mathsf{pos} \in [2:7]$.

Sparse representation. SSTs are tailored toward representing sparse arrays. This allows the height of the tree to shrink, which consequently results in more efficient updates and queries. A crucial aspect of SSTs is that they are self-adapting and do not require any a priori information about the sparsity of the input.

Intuitively, we achieve this as follows. Empty array indexes are never explicitly represented in the intermediate nodes of the tree. Every intermediate node nd is such that nd.pos points to a unique, non-empty entry of A. The opposite is also true: every non-empty entry of A is indexed by the pos field of a node in the earliest possible level satisfying Equation (2).

Example 4. Consider an array A of length 8, initially empty. Figure 6(a)– (d) sequentially updates A with new entries, and Figure 6(e)– (h) shows the evolution of the corresponding SST T. In Figure 6(e), the unique node of T represents that A has only one non-empty entry, at index 2 with value 65. In contrast, a standard (non-sparse) ST would contain 15 nodes and have height 3. Figure 6(f) shows the update of T after setting $A[3] = 42$. Since 42 is now the smallest value of A, it is stored in the root, while a new node $\mathsf{nd}_{2,2}$ becomes the left child of the root, with its pos and min fields set according to Equation (2). Again, in the standard ST representation, $\mathsf{nd}_{2,2}$ would have nodes $\mathsf{nd}_{0,3}$ and $\mathsf{nd}_{2,3}$ as ancestors, but these nodes are never created in our SST. Figure 6(g) shows another update of T after setting $A[0] = 59$. As $59 \leq \mathsf{nd}_{0,7}.\mathtt{min}$, the root remains intact. The index 0 belongs to the left subtree of the root and $59 < \mathsf{nd}_{2,2}.\mathtt{min}$. This causes the new entry to be inserted to a fresh node $\mathsf{nd}_{0,3}$ and assigned as the left child of the root, taking the existing node $\mathsf{nd}_{2,2}$ as its right child. Figure 6(h) shows the final update of T after setting $A[7] = 13$. Since 13 is now the smallest value of A, it is stored in the root, while the previous value stored in the root is recursively pushed downwards.

One last optimization step is to flatten subtrees of small size at the leaves to array blocks. For a node $\mathsf{nd}_{i,j}$, if $\mathtt{depth}(\mathsf{nd}_{i,j})$ is larger than a value deeming that the node represents a small sub-array

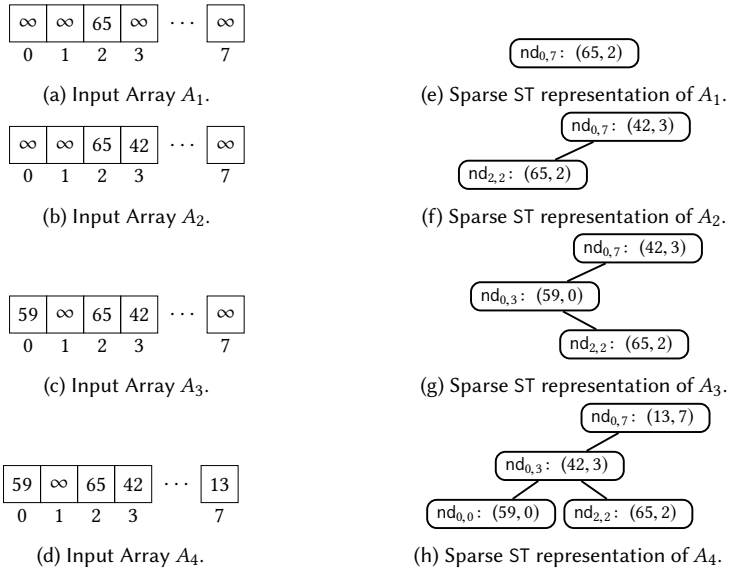

Fig. 6. Sparse ST representation. The notation $\mathsf{nd}_{i,j} : (x, y)$ denotes $\mathsf{nd}_{i,j}.\mathtt{min} = x$, $\mathsf{nd}_{i,j}.\mathtt{pos} = y$.

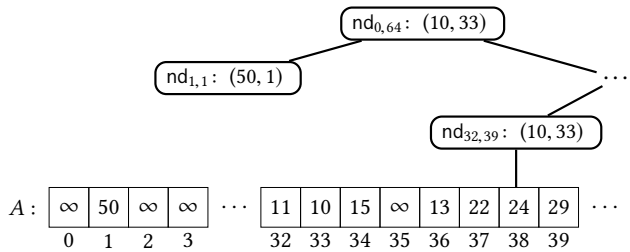

Fig. 7. Flatting small subtrees to block nodes.

of A, then it becomes a leaf in T, storing the subarray $A[i : j]$. We call such a node $\mathsf{nd}_{i,j}$ a *block node*. This representation is valuable when the array contains segments which are densely populated but also far apart from each other.

Example 5. Consider the array A and its corresponding SST in Figure 7. A contains only a single entry in the range $0, \ldots, 31$. This entry is sparsely represented in the left subtree of the root. On the other hand, A is dense in the range $32, \ldots, 39$, which would require a full subtree. Instead, when the range of the represented subarray is smaller than a threshold b, the corresponding subtree is flattened to an array block.

Operations on Sparse Segment Trees. Algorithm 1 implements the operations `update`, `min`, `argleq` on an SST.

Function `update`(nd, *pos*, *v*). This function inserts a node nd′ into the subtree of nd, with $\mathsf{nd}'.\mathtt{pos} = pos$ and $\mathsf{nd}'.\mathtt{min} = v$. It first checks if nd is initialized; if not, it creates a new node (Line 3). If nd is the root, the function checks and deletes if a node nd″ exists in the subtree of nd where $\mathsf{nd}''.\mathtt{pos} = pos$ (Line 5). Then, it checks if $2^{\mathsf{depth}(\mathsf{nd})} \geq n/b$, where b is the block-size threshold. If so, nd becomes a block node and the update is performed directly on nd.`block` (Line 7). Otherwise,

ALGORITHM 1: The Sparse Segment Tree.

```
1  function update(nd, pos, v)
2      if nd = nil then
3          nd ← createNode(pos, v)
4      else if nd = T.root then
5          nd ← deleteNode(nd, pos)
6      else if 2^depth(nd) ≥ n/b then
7          nd.block[pos] = v;
8      else
9          if v < nd.min ∨ (v = nd.min ∧ pos > nd.pos) then
10             nd.min, nd.pos, v, pos ← v, pos, nd.min, nd.pos
11         Let mid ← nd.start + ⌊(nd.end − nd.start)/2⌋
12         if pos ≤ mid then
13             nd.left ← updateHelper(nd.left, pos, v)
14         else
15             nd.right ← updateHelper(nd.right, pos, v)
16 function updateHelper(nd, pos, v)
17     if nd = nil then
18         nd ← createNode(pos, v)
19     else if pos ∈ [nd.start, nd.end] then
20         update(nd, pos, v)
21     else
22         nd ← createLowestCommonAncestor(nd, pos, v)
23     return nd
24 function min(nd, i)
25     if nd = nil ∨ i > nd.end then
26         return ∞
27     if nd.block ≠ nil then
28         return min(nd.block, i);
29     if nd.pos ≥ i then
30         return nd.min;
31     else
32         l ← min(nd.left, i)
33         r ← min(nd.right, i)
34         if l < r then
35             return l
36         else
37             return r
38 function argleq(nd, v)
39     if nd = nil ∨ nd.min > v then
40         return −∞
41     if nd.block ≠ nil then
42         return argleq(nd.block, v);
43     if nd.pos ≥ nd.left.end ∧ nd.pos ≥ nd.right.end
       then
44         return nd.pos
45     else if nd.right.min ≤ v then
46         return max(nd.pos, argleq(nd.right, v))
47     else
48         return max(nd.pos, argleq(nd.left, v))
```

it checks if $v \leq$ nd.min. If so, it swaps the values of nd.pos and nd.min with *pos* and v (Line 10). This captures that the new entry will be stored in nd and the previous entry in nd will be relocated to a child node. Next, it identifies which subtree the corresponding entry belongs to, and makes a call to updateHelper (Lines 13, 15). If the subtree is not present, updateHelper creates a new node (Line 18). Otherwise, it checks if pos is within the range of the subtree and if so it recursively calls update. If the subtree is present but pos is not within the range, createLowestCommonAncestor creates a new node whose range corresponds to the lowest common ancestor of the existing subtree and pos.

Function min(nd, i). This recursive function returns ∞ if nd is nil or $i >$ nd.end (Line 26), which captures the end of the recursion. Otherwise, it returns the smallest min value in the subtree of nd whose corresponding pos is greater than or equal to i, as follows. First, it checks if nd is a block node, in which case it returns the minimum of the block array nd.block (Line 28). Otherwise, if the minima indexing optimization succeeds, the traversal stops here, returning nd.min (Line 30). If both conditions fail, the function proceeds recursively on the left and right child of nd (Lines 31– 33).

Function argleq(nd, v). This function returns $-\infty$ if nd is nil or nd.min $> v$. Otherwise, it returns the largest pos in the subtree of nd whose corresponding min value is smaller than or equal to v, as follows. It first checks if nd is a block node, in which case it simply performs an array traversal on nd.block (Line 42). Otherwise, the function uses the minima indexing information in nd.pos to

check if it can return early (Line 44). If the check fails, it makes a recursive call on the right subtree if $\texttt{nd.right.min} \leq v$ (Line 46). If all conditions fail, it makes a recursive call on the left subtree (Line 48).

Using SSTs. SSTs handle operations on the represented suffix-minima array A is as follows. Queries $\texttt{min}(A, i)$ and $\texttt{argleq}(A, a)$ are handled as $\texttt{min}(T.\texttt{root}, i)$ and $\texttt{argleq}(T.\texttt{root}, a)$, respectively. On the other hand, $\texttt{update}(A, i, a)$ is handled as $\texttt{update}(T.\texttt{root}, i, a)$.

Complexity and the height of SSTs. The theoretical advantage of SSTs over STs is the fact that their height is bounded not only by $\log n$, but also by the density of the represented array. This is captured in the following lemma.

Lemma 1. *Consider an array A of size n with d non-empty (i.e., non-∞) entries, and let T be its corresponding* SST. *Then the height of T is bounded by* $\min(\log n, d)$.

Proof. Observe that every intermediate node $\texttt{nd}_{i,j}$ of an SST T has its field $\texttt{nd}_{i,j}.\texttt{pos}$ point to a non-empty entry of $A[i, j]$. Due to Equation (2), this is a different entry from those pointed by the pos fields of the ancestors of $\texttt{nd}_{i,j}$. Moreover, any node $\texttt{nd}_{i',j'}$ which is neither an ancestor nor a descendant of $\texttt{nd}_{i,j}$ has a non-overlapping range with $\texttt{nd}_{i,j}$, i.e., $[i : j] \cap [i' : j'] = \emptyset$. It follows that $\texttt{nd}_{i,j}.\texttt{pos}$ indexes a unique entry of A. In turn this implies that the height of T is bounded by the number d of non-empty elements of A.

In addition, as in standard STs, for every two nodes $\texttt{nd}_{i,j}$ and $\texttt{nd}_{i',j'}$ such that $\texttt{nd}_{i,j}$ is the parent of $\texttt{nd}_{i',j'}$, we have that $(j' - i') \leq \lceil (j - i)/2 \rceil$, i.e., the range halves in each step. Thus the height T is also bounded by $\log n$, as desired. □

Thus, we can use Lemma 1 to bound the time taken to perform any update and query on the represented array A. In practice, however, the running time can be even smaller, as, in contrast to standard STs, the update and query operations can stop before traversing a full path from the root to a leaf.

3.3 Collective Sparse Segment Trees

We are now ready to describe CSSTs, leading to Theorem 1.

CSSTs. For every two distinct chains $t_1, t_2 \in [k]$, CSSTs maintain a suffix-minima array $\texttt{A}_{t_1}^{t_2}$ storing direct edges from t_1 to t_2, together with a min-heap $\texttt{edgeHeap}_{t_1}^{t_2}[j_1]$, for each $j_1 \in [n]$, as described in Section 3.1. Suffix-minima queries on $\texttt{A}_{t_1}^{t_2}$ are handled using an SST, as described in Section 3.2. The pseudocode is shown in Algorithm 2. The functions `insertEdge` and `deleteEdge` follow directly the description in Section 3.1, i.e., they manipulate the arrays $\texttt{A}_{t_1}^{t_2}$ and min-heaps $\texttt{edgeHeap}_{t_1}^{t_2}[j_1]$.

The function `reachable` is implemented via a `successor` query, which is more involved. As we have $k > 2$ chains, `successor` runs a forward traversal from $\langle t_1, j_1 \rangle$ to discover transitive reachability across chains. This is done by repeated `min` queries in the suffix minima arrays $\texttt{A}_{t_2'}^{t_1'}$ (Line 9), which find the earliest node of t_1' that has an edge from the currently earliest node of t_2' found to be reachable from $\langle t_1, j_1 \rangle$. Crucially, this computation completes in $O(k^3)$ steps, as opposed to $O(n^2)$ steps normally required for a graph traversal.

Example 6. Consider the chain DAG in Figure 8 and the operation $\texttt{successor}(\langle 0, 0 \rangle, 3)$. The first step is identifying the earliest successors of $\langle 0, 0 \rangle$ reachable via direct edges from $\langle 0, 0 \rangle$ and its thread-local successors, in all other threads. This discovers $\langle 1, 0 \rangle$ via (1) and $\langle 3, 2 \rangle$ via (2). While $\langle 3, 2 \rangle$ is indeed in chain 3, it is not necessarily the *earliest* successor of $\langle 0, 0 \rangle$ (indeed, the earliest successor is $\langle 3, 1 \rangle$). Hence the function proceeds iteratively for each newly discovered node until reaching a fixed point. This leads to the discovery of $\langle 2, 1 \rangle$ via (3) and eventually the discovery of $\langle 3, 1 \rangle$ via (4).

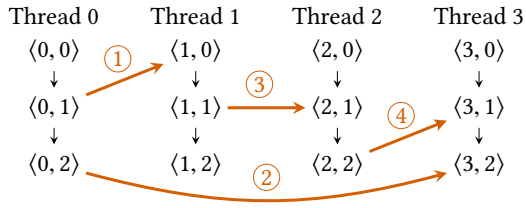

Fig. 8. Query successor($\langle 0,0\rangle, 3$).

ALGORITHM 2: Fully dynamic CSSTs.

```
1  function successor(⟨t1, j1⟩, t2)
2      foreach t1' ∈ [k] \ {t1} do
3          [t1'] ← min(A_{t1}^{t1'}, j1)
4      do
5          changed ← false
6          foreach t1' ∈ [k] \ {t1} do
7              foreach t2' ∈ [k] \ {t1} do
8                  if t1' ≠ t2' then
9                      Let v ← min(A_{t2'}^{t1'}, [t2'])
10                     if v < [t1'] then
11                         [t1'] ← v
12                         changed ← true
13     while changed
14     return [t2]

15 function reachable(⟨t1, j1⟩, ⟨t2, j2⟩)
16     if successor(⟨t1, j1⟩, t2) ≤ j2 then
17         return True
18     else
19         return False

20 function insertEdge(⟨t1, j1⟩, ⟨t2, j2⟩)
21     if j2 < edgeHeap_{t1}^{t2}[j1].min then update(A_{t1}^{t2}, j1, j2)
22     edgeHeap_{t1}^{t2}[j1].insert(j2)

23 function predecessor(⟨t1, j1⟩, t2)
24     foreach t1' ∈ [k] \ {t1} do
25         [t1'] ← argleq(A_{t1'}^{t1}, j1)
26     do
27         changed ← false
28         foreach t1' ∈ [k] \ {t1} do
29             foreach t2' ∈ [k] \ {t1} do
30                 if t1' ≠ t2' then
31                     Let v ← argleq(A_{t1'}^{t2'}, [t2'])
32                     if v > [t1'] then
33                         [t1'] ← v
34                         changed ← true
35     while changed
36     return [t2]

37 function deleteEdge(⟨t1, j1⟩, ⟨t2, j2⟩)
38     if j2 = edgeHeap_{t1}^{t2}[j1].min then
39         edgeHeap_{t1}^{t2}[j1].delete(j2)
40         update(A_{t1}^{t2}, j1, edgeHeap_{t1}^{t2}[j1].min)
41     else
42         edgeHeap_{t1}^{t2}[j1].delete(j2)
```

Analysis. We now state the properties of fully dynamic CSSTs. First, since the arrays $\mathtt{A}_{t_1}^{t_2}$ only store direct edges, the following sparsity property is straightforward.

Lemma 2 (Sparsity). *Let d be the cross-chain density of G. Then the density of each array $\mathtt{A}_{t_1}^{t_2}$ is bounded by d.*

Proof. Since we have $\mathtt{A}_{t_1}^{t_2}[j_1] \neq \infty$ iff there exists some $j_2 \in [n]$ with $\langle t_1, j_1\rangle \to \langle t_2, j_2\rangle$, it follows that the density of $\mathtt{A}_{t_1}^{t_2}$ is bounded by the cross-chain density d of G. □

The use of min-heaps guarantees the following invariant, which informally states that each SST stores the earliest outgoing neighbor of each node to each respective chain.

Lemma 3 (Invariant). *For every distinct $t_1, t_2 \in [k]$, we have $\mathsf{A}^{t_2}_{t_1}[j_1] = \min(\{j_2 \in [n].\ \langle t_1, j_1\rangle \to \langle t_2, j_2\rangle\})$.*

Proof. The statement follows from the invariant that $\mathsf{A}^{t_2}_{t_1}[j_1]$ is equal to the root of the min-heap $\mathtt{edgeHeap}^{t_2}_{t_1}[j_1]$, which, at all times, contains all $j_2 \in [n]$ such that $\langle t_1, j_1\rangle \to \langle t_2, j_2\rangle$. □

A *crossing path* is a sequence of nodes $\pi : \langle t_0, j_0\rangle, \ldots, \langle t_{m-1}, j_{m-1}\rangle$ such that for each $\ell \in [m-2]$ we have $t_\ell \neq t_{\ell+1}$ and there are $i_1, i_2 \in [n]$ such that (i) $i_1 \geq j_\ell$, (ii) $i_2 \leq j_{\ell+1}$, and (iii) $\langle t_\ell, i_1\rangle \to \langle t_{\ell+1}, i_2\rangle$. Clearly, for any two nodes $\langle t_1, j_1\rangle$ and $\langle t_2, j_2\rangle$, if $\langle t_1, j_1\rangle \to^* \langle t_2, j_2\rangle$, then there exists a crossing path between them of length at most k. The following lemma implies the correctness of queries.

Lemma 4. *For every $i \in [k]$, after the i-th iteration of the do-while loop in Line 4, $[t_1']$ points to the earliest node in chain t_1' that is reachable from $\langle t_1, j_1\rangle$ via a crossing path of length $\leq i+1$.*

Proof. The statement follows by induction on i, in a style similar to correctness proof of the Bellman–Ford algorithm. In particular, for $i = 0$, every crossing path consists of two nodes $\pi : \langle t_1, j_1\rangle, \langle t_2, j_2\rangle$, and indeed, $[j_1] = j_1$, by initialization loop in Line 2.

Now assume that the statement holds for some i, and we argue that it holds for $i+1$. Consider any crossing path $\pi : \langle t_1, j_1\rangle, \ldots, \langle t_2, j_2\rangle$ of length $i+1$, and let π' be its prefix of length i, ending in some node $\langle t_3, j_3\rangle$. By the induction hypothesis, we have that after the ith iteration of the do-while loop, we have $[j_1] \leq j_3$. Moreover, since the last pair of π is $\ldots \langle t_3, j_3\rangle, \langle t_2, j_2\rangle$, there exist $i_1, i_2 \in [k]$ with $i_1 \geq j_3$ and $i_2 \leq j_2$ such that $\langle t_3, i_1\rangle \to \langle t_2, i_2\rangle$. Then the suffix minima query in Line 9 will obtain $v \leq i_2$, and thus the invariant is restored in Line 11. □

Regarding the time complexity, due to Lemma 2 and 1, each operation in a suffix minima array $\mathsf{A}^{t_2}_{t_1}$ takes $O(\min(\log n, d))$ time. Since the maximum out-degree is δ, each operation on a min heap $\mathtt{edgeHeap}^{t_2}_{t_1}[j_1]$ runs in $O(\log \delta)$ time. Thus, updates take $O(\max(\log \delta, \min(\log n, d)))$ time, while, due to Lemma 4, queries take $O(k^3 \min(\log n, d))$ time, arriving at Theorem 1.

Space usage. We now discuss the space usage of CSSTs. Consider a chain DAG G of n nodes, k chains, and cross-chain density d. In this fully dynamic setting, CSSTs store all edges in the min heaps $\mathtt{edgeHeap}^{t_2}_{t_1}[j_1]$, as edge deletions requires to remember all edges added so far. The suffix minima arrays $\mathsf{A}^{t_2}_{t_1}$ store in total $O(dk)$ entries, as (i) each chain has at most d entries that have at least one successor to another chain, and (ii) there are $< k$ other chains. Naturally, $d \leq n$, thus in the worst case the above expression becomes $O(nk)$. This is the same space complexity as Vector Clocks. However, Vector Clocks do not exploit the sparsity of G. As we will see in Section 5, in practice typically $d \ll n$, which makes CSSTs more space-efficient than Vector Clocks.

Thread safety. Thread safety becomes a consideration when a graph must be maintained dynamically by several threads, for example when analyzing concurrent executions in an online monitoring setting. Such interference may also arise when using CSSTs, and can lead to incorrect outcomes unless properly synchronized. A straightforward strategy for synchronization is to use a global read/write lock, acquiring a read lock for queries `reachable`, `successor`, and `predecessor`, and a write lock for update operations `insertEdge` and `deleteEdge`. We leave the design of more fine-grained synchronization schemes to future work.

4 Incremental CSSTs

In practice, many dynamic analyses maintain their partial order incrementally, i.e., they only insert and never delete orderings. To optimize for this incremental setting, in this section we specialize CSSTs to only handle incremental updates, with guarantees stated in the following theorem.

ALGORITHM 3: Incremental CSSTs.

```
1  function successor(⟨t1, j1⟩, t2)
2      return min(A^{t2}_{t1}, j1)
3  function predecessor(⟨t1, j1⟩, t2)
4      return argleq(A^{t1}_{t2}, j1)
5  function reachable(⟨t1, j1⟩, ⟨t2, j2⟩)
6      if successor(⟨t1, j1⟩, t2) ≤ j2 then
7          return True
8      else
9          return False
10 function insertEdge(⟨t1, j1⟩, ⟨t2, j2⟩)
11     foreach t'1 ∈ [k] do
12         foreach t'2 ∈ [k] \ {t'1} do
13             if t'1 = t1 then
14                 Let j'1 ← ⟨t1, j1⟩
15             else
16                 Let j'1 ← predecessor(⟨t1, j1⟩, t'1)
17             if t'2 = t2 then
18                 Let j'2 ← ⟨t2, j2⟩
19             else
20                 Let j'2 ← successor(⟨t2, j2⟩, t'2)
21             if successor(⟨t'1, j'1⟩, t'2) > j'2 then
22                 update(A^{t'2}_{t'1}, j'1, j'2)
```

Theorem 2. *Consider a chain DAG G of n nodes, k chains, and cross-chain density d. Incremental* CSSTs *maintain G under incremental updates, costing $O(k^2 \min(\log n, d))$ time per update and $O(\min(\log n, d))$ time per query.*

Compared to Theorems 1, 2 trades the cost dependency on the number of chains from queries to updates, while also decreasing it by a factor k. Moreover, now the cost of both queries and updates is bounded by the cross-chain density of G (as opposed to only queries in Theorem 1).

Incremental CSSTs**.** Incremental CSSTs maintain $k(k-1)$ suffix minima arrays $\mathtt{A}^{t_2}_{t_1}$, with $t_1 \neq t_2$, each implemented using an SST $T^{t_2}_{t_1}$. These arrays now store *transitive reachability*, as opposed to direct edges between chains. Algorithm 3 shows how each array $\mathtt{A}^{t_2}_{t_1}$ is queried and updated, following the reachability queries and edge insertions in CSSTs. The functions successor$(\langle t_1, j_1\rangle, t_2)$, predecessor$(\langle t_1, j_1\rangle\rangle, t_2)$, and reachable$(\langle t_1, j_1\rangle, \langle t_2, j_2\rangle)$ handle successor, predecessor, and reachability queries in a straightforward manner: since each $\mathtt{A}^{t_2}_{t_1}$ is transitively closed, these operations reduce to suffix-minima queries on $\mathtt{A}^{t_2}_{t_1}$.

The function insertEdge$(\langle t_1, j_1\rangle, \langle t_2, j_2\rangle)$ is somewhat more complex, to guarantee that the arrays $\mathtt{A}^{t_2}_{t_1}$ indeed store transitive reachability across chains. To achieve this, besides inserting the direct edge $\langle t_1, j_1\rangle \to \langle t_2, j_2\rangle$, this function performs a transitive closure computation by finding the predecessor $\langle t'_1, j'_1\rangle$ of $\langle t_1, j_1\rangle$ (Lines 14– 16) and the successor $\langle t'_2, j'_2\rangle$ of $\langle t_2, j_2\rangle$ (Lines 18– 20) in all pairs of chains t'_1, t'_2. Then, an edge $\langle t'_1, j'_1\rangle \to \langle t'_2, j'_2\rangle$ is inserted unless a path between these nodes already exists.

Example 7. Consider the chain DAG in Figure 9(a) and the operation insertEdge$(\langle 1, 1\rangle, \langle 2, 0\rangle)$. We obtain the transitive path $\langle 0, 1\rangle \to^* \langle 3, 2\rangle$ as follows. First, the predecessor of $\langle 1, 1\rangle$ in thread 0 is identified via the query predecessor$(\langle 1, 1\rangle, 0)$. This query performs an argleq operation on $\mathtt{A}^1_0$ (Figure 9(b)) with value 1, returning node $\langle 0, 1\rangle$. Then, the successor of $\langle 2, 0\rangle$ in thread 3 is identified via the query successor$(\langle 2, 0\rangle, 3)$. This query performs a min operation on $\mathtt{A}^3_2$ (Figure 9(c)) with index 0, and returns the node $\langle 3, 2\rangle$. Next, the update operation on $\mathtt{A}^3_0$ adds the transitive edge, resulting in $\mathtt{A}^3_0$ in Figure 9(d).

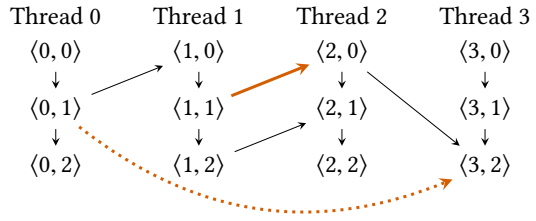

(a) A chain DAG G.

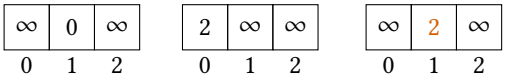

(b) Array $\mathtt{A}^1_0$. (c) Array $\mathtt{A}^3_2$. (d) Array $\mathtt{A}^3_0$.

Fig. 9. Inserting an edge $\langle 1,1\rangle \to \langle 2,0\rangle$ and inferring the transitive path $\langle 0,1\rangle \to^* \langle 3,2\rangle$.

Analysis. The correctness of incremental CSSTs is based on the following invariants. The first concerns soundness, i.e., every entry in each array $\mathtt{A}^{t_2}_{t_1}$ corresponds to a path between the respective nodes in chains t_1 and t_2. The second concerns completeness, i.e., if there is a path $\langle t_1, j_1\rangle \to^* \langle t_2, j_2\rangle$ between two nodes, it is captured as a suffix minima in the array $\mathtt{A}^{t_2}_{t_1}$. In conjunction, they imply the correctness of `successor`, `predecessor`, and `reachable` queries implemented as `min` and `argleq` operations on the corresponding array $\mathtt{A}^{t_2}_{t_1}$.

Lemma 5 (Soundness). $\mathtt{A}^{t_2}_{t_1}[j_1] = j_2 \implies \langle t_1, j_1\rangle \to^* \langle t_2, j_2\rangle$.

Proof. Consider any update $\mathtt{update}(\mathtt{A}^{t'_2}_{t'_1}, j'_1, j'_2)$ happening in Line 22 when the algorithm processes a call $\mathtt{insertEdge}(\langle t_1, j_1\rangle, \langle t_2, j_2\rangle)$, and we will argue that $\langle t'_1, j'_1\rangle \to^* \langle t'_2, j'_2\rangle$. Indeed, inductively, we have $\langle t'_1, j'_1\rangle \to^* \langle t_1, j_1\rangle$ and $\langle t_2, j_2\rangle \to^* \langle t'_2, j'_2\rangle$, thus after the end $\langle t_1, j_1\rangle \to \langle t_2, j_2\rangle$ is inserted in G, we have $\langle t'_1, j'_1\rangle \to^* \langle t'_2, j'_2\rangle$, as desired. □

Lemma 6 (Completeness). $\langle t_1, j_1\rangle \to^* \langle t_2, j_2\rangle$ *and* $t_1 \neq t_2 \implies \exists j'_1 \geq j_1.\ \mathtt{A}^{t_2}_{t_1}[j'_1] \leq j_2$.

Proof. Consider any two nodes $\langle t'_1, j'_1\rangle \to^* \langle t'_2, j'_2\rangle$ after the algorithm processes a call $\mathtt{insertEdge}(\langle t_1, j_1\rangle, \langle t_2, j_2\rangle)$, and we will argue that there exists some $i'_2 \in [n]$ such that $i'_2 \leq j'_2$ and $\mathtt{A}^{t'_2}_{t'_1}[j_1] = i'_2$. The statement holds by induction if there is a path $\langle t'_1, j'_1\rangle \to^* \langle t'_2, j'_2\rangle$ not containing the edge $\langle t_1, j_1\rangle \to \langle t_2, j_2\rangle$. Otherwise, we have (i) a path $\langle t_2, j_2\rangle \to^* \langle t'_2, j'_2\rangle$, and (ii) a path $\langle t'_1, j'_1\rangle \to^* \langle t'_1, j'_1\rangle$. These two paths are merged appropriately in Line 22, leading to $\mathtt{A}^{t'_2}_{t'_1}[j_1] = i'_2$, as desired. □

Recall that the key feature of SSTs is their ability to handle sparse arrays efficiently. As incremental CSSTs insert transitive edges in the arrays $\mathtt{A}^{t_2}_{t_1}$, this might increase their density, adversely affecting the underlying SST, as we saw in Section 3.2. At closer inspection, however, the density of each $\mathtt{A}^{t_2}_{t_1}$ cannot grow beyond the cross-chain density of G.

Lemma 7 (Sparsity). *Let d be the cross-chain density of G. Then the density of each array $\mathtt{A}^{t_2}_{t_1}$ is bounded by d.*

Proof. The statement follows from the following observation. Whenever Algorithm 3 inserts an edge $\mathtt{update}(\mathtt{A}^{t'_2}_{t'_1}, j'_1, j'_2)$, we have that either (i) $t'_1 = t_1$, in which case $j'_1 = j_1$, or (ii) $j_1 =$

predecessor($\langle t_1, j_1\rangle, t_1'$). In the first case we are inserting an outgoing edge to $\langle t_1', j_1'\rangle$, while in the second case, $\langle t_1', j_1'\rangle$ already has an edge to some other chain. In both cases, $\langle t_1', j_1'\rangle$ counts toward the cross-chain density d of G. Hence, the number of non-∞ entries of each suffix-minima array is bounded by d, as desired. □

Combined with our bounds on the height of SSTs (Lemma 1 in Section 3.2), Lemma 7 implies that the height of the SST representing $\mathtt{A}_{t_1}^{t_2}$ is bounded both by $\log n$ and by the cross-chain density of G. We thus arrive at Theorem 2.

Space usage. Due to Lemma 7, the total space used by the suffix minima arrays $\mathtt{A}_{t_1}^{t_2}$ is $O(dk)$, by an analysis similar to that of fully dynamic CSSTs in Section 3.3. This time, CSSTs do not use heaps to store all edges, thus the total space used is $O(dk)$. Vector Clocks and STs [32], on the other hand, use $O(nk)$ space, which is larger in practice.

5 Experimental Evaluation

Here we incorporate CSSTs in a number of dynamic analyses from recent literature and evaluate their performance.

5.1 Experimental Setup

Data structures. Each analysis dynamically maintains a partial order P using various representations. We explore different standard ways for this purpose, as follows.

— Graphs is a standard graph representation of P, that is not transitively closed.
— VCs is a standard Vector Clock representation of P, that is (by design) transitively closed [29].
— STs is the data structure in [32] based on STs.
— CSSTs, as introduced in this article.

Most analyses make only incremental updates, in which case Vector Clocks are preferred over plain graphs. Thus in this setting we compare CSSTs to VCs and ST. On the other hand, VCs and ST cannot handle decremental updates, hence in the fully dynamic setting we compare CSSTs to Graphs. The data structure used in the respective paper is marked by †. As some tools are not public, we implemented their analysis following the corresponding paper.

Optimizations. To set the threshold b for the size of the block nodes in CSSTs (Section 3.2), we perform a randomized stress test with varying sizes of b, and measure their performance. Based on this test, we set $b = 32$.

We also implement two optimizations in VCs. First, each new edge $e_1 \to e_2$ must be propagated to the successors of e_2 in its own chain. We stop this propagation as soon as we find some e_2' with already $e_1 \to^* e_2'$. Second, for each chain, we do not create VCs for the suffix of its events that do not have any direct orderings from other chains, as ordering queries on such nodes can be inferred from their predecessors.

Setup. In each analysis, we use the dataset of the corresponding paper and relevant literature, and report on benchmarks on which some data structure runs in $\geq$ 1s. For some bigger datasets, some of the data structures lead to a timeout or memoryout. In each of these cases, we add the timeout time to the total time taken for processing the whole benchmark. We use an Ubuntu 22.04 machine with 2.4 GHz CPU and 64 GB of memory. We repeat each experiment 10 times and report the average time/memory measurement.

5.2 Experimental Results

We begin with a macroscopic view in Figures 10 and 11, showing the (geometric) average improvement in time and memory that CSSTs offer over VCs, STs and Graphs in each analysis. We see that in

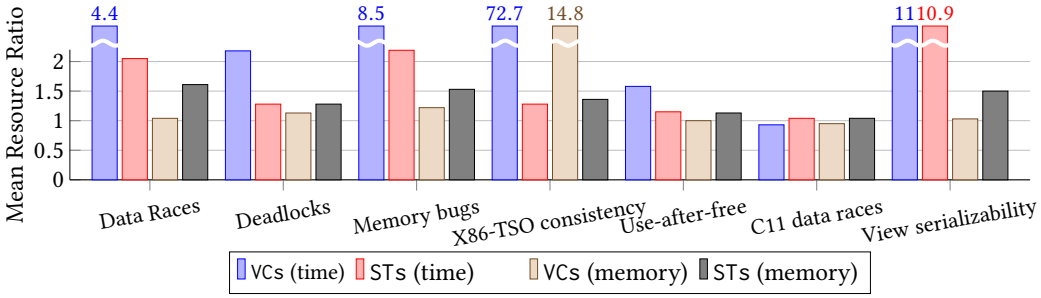

Fig. 10. The geometric mean of the ratio of time and memory used in each benchmark by VCs and STs over incremental CSSTs in the dataset of the corresponding analysis.

nearly all cases, these ratios are above 1, meaning that CSSTs indeed outperform the other data structures. The only exception concerns data races in C11; we will analyze this case in more detail later. The improvement is even more pronounced on the last analysis (linearizability), which is fully dynamic and thus the only baseline is plain Graphs.

Although the worst-case memory usage of CSSTs is $O(nk)$, and thus the same as VCs and STs, in practice it is less. This confirms our insights that typical analyses give rise to sparse chain DAGs, which is then exploited by CSSTs. We will take a closer look into the sparsity of each benchmark later.

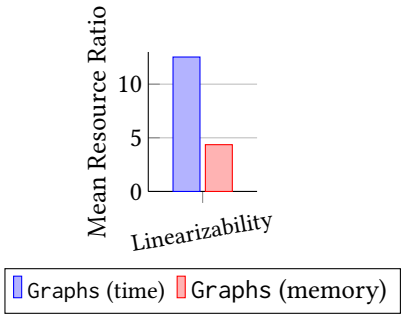

Fig. 11. The geometric mean of the ratio of time and memory for the Linearizability benchmark.

We remark that the resource (time, memory) reported in each benchmark concerns the whole analysis, not just the corresponding data structure.[1] The actual resource improvement achieved by CSSTs is significantly larger.

1. Data race prediction. We start with the M2 data race detector [32]. This analysis observes traces σ which might be data-race free, and attempts to permute them to new traces σ' that are valid (called correct reorderings) and expose a data race. Figure 12 shows such an example. Internally, M2 utilizes STs for maintaining its partial order P. The reachability queries are incremental. We measure the total time taken in the analysis.

Table 2 shows the performance of M2 with each data structure. VCs become unscalable early, and are outperformed by STs, which were indeed developed to address this scalability issue. We observe that CSSTs are the most performant data structure on each individual benchmark. Their advantage over STs is primarily due to the efficient way that CSSTs handle sparsity (Section 3.2). Indeed, column q reports the mean density among each suffix minima array inside CSSTs when it obtained its densest form, and matches our expectations that is typically small. Our sparse treatment not only reduces the runtime of CSSTs but also their memory footprint, and thus support the analysis even when VCs and STs go out of memory (on xalan and batik).

As a side note, observe that the running time of each method is not always increasing when we move to larger traces (e.g., all data structures take more time to process moldyn, with size $N = 200.3$K than jigsaw, with size 3.1M). This is because moldyn has more potential data races

[1]The analysis on data races for C11 is an exception, as we explain later.

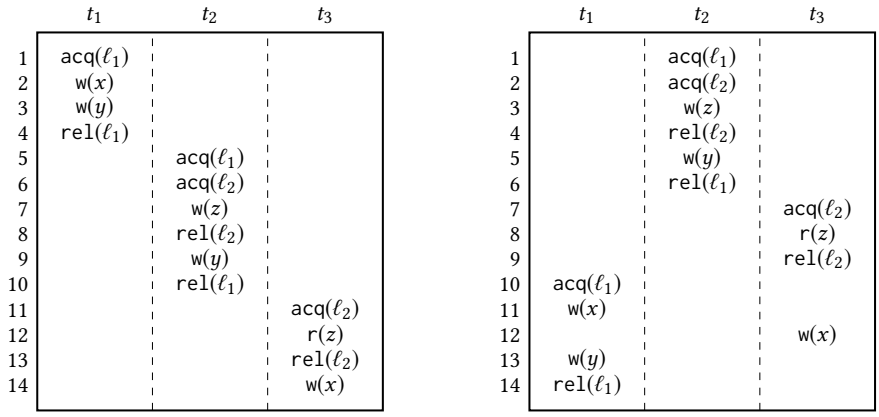

(a) In trace σ, (e_2, e_{14}) is not a race.

(b) In trace σ', (e_{11}, e_{12}) forms a race.

Fig. 12. An example analysis of M2, taken from [32]. On the left, a data-race-free trace σ is processed. On the right, M2 infers a valid permutation σ' of σ that exposes a data race, between the two writes on x. Intuitively, every read event observes the same write event in both σ and σ', which in turn guarantees that σ' is a valid trace of any program that generated σ.

Table 2. Race Prediction Results

benchmark	T	N	q	VCs (s)	STs† (s)	CSSTs (s)
clean	12	1.3K	0.28	2.4	2.6	1.6
bubblesort	29	4.7K	0.15	69.4	56.2	27.5
lang	10	6.3K	0.22	93.3	39.2	27.3
readerswriters	8	11.3K	0.32	54.4	23.5	18.5
raytracer	6	15.8K	0.17	10.5	10.2	9.3
bufwriter	9	22.3K	0.2	130	24.6	12.9
ftpserver	14	49.6K	0.06	41.2	14.1	9.0
moldyn	6	200.3K	0.12	T.O.	T.O.	12636
linkedlist	15	1.0M	0.06	T.O.	T.O.	8893
derby	7	1.4M	0.04	1436	215	197
jigsaw	15	3.1M	0.06	154	45.6	32.0
sunflow	17	11.7M	0.01	T.O.	780	505
xalan	9	122.5M	0.01	T.O.	O.O.M.	979
batik	8	157.9M	0.01	O.O.M.	O.O.M.	1956
Total	-	**297.9M**	-	**>91992**	**>73211**	**25304**

The timeout is 5h. Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

that the analysis needs to check, despite its smaller size. Similar, non-monotonic, patterns appear in other analyses as well.

2. Deadlock prediction. Here we incorporate CSSTs in SeqCheck, a dynamic deadlock prediction analysis [9]. This analysis identifies potential deadlocks by analyzing lock-acquisition orders between threads, and then try to witness each deadlock by a valid reordering of the observed trace. In Figure 13, we show an example of a trace σ that doesn't witness a deadlock, however, (e_1, e_2) and (e_9, e_{10}) may incur a deadlock potentially. For this purpose, it incrementaly maintains a partial

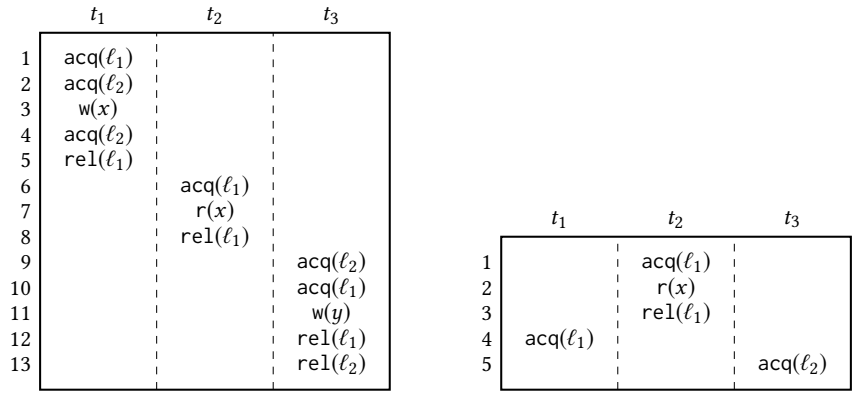

(a) Trace σ has no deadlock. (b) Trace σ' has a deadlock.

Fig. 13. An example of the deadlock, as discussed in [9], involves trace σ, which includes the events e_1, e_2, e_9, and e_{10}. If the program that generated σ be executed in the order of σ', a deadlock occurs because t_1 holds ℓ_1 and waits for ℓ_2, while t_3 holds ℓ_2 and waits for ℓ_1.

Table 3. Deadlock Results

benchmark	T	N	q	VCs (s)	STs† (s)	CSSTs (s)
jigsaw	21	143.0K	0.03	356	14.8	7.8
elevator	5	245.9K	0.02	3.7	2.8	2.6
hedc	7	409.8K	0.04	23.2	22.0	23.2
JDBCMySQL	3	442.9K	0.01	3.7	4.1	3.5
cache4j	2	775.5K	0.01	5.5	6.6	5.6
Swing	8	3.8M	0.01	14.8	13.4	13.1
sunflow	15	21.5M	0.01	369	111	110
eclipse	15	64.2M	0.01	535	711	282
Total	**-**	**91.4M**	**-**	**1311**	**886**	**448**

Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

order which will eventually be linearized to this valid reordering. To maintain the partial order efficiently, SeqCheck utilizes STs internally.

Our results are shown in Table 3. VCs are again the slowest performer by far, though they are very competitive on a few benchmarks for which deadlocks can be detected fairly early in the analysis. Naturally, STs and CSSTs perform similarly on these benchmarks as well. On the more demanding benchmarks, however, CSSTs are clearly the best performer. We also observe that the average density q is small.

3. Memory bug prediction. Here we incorporate CSSTs in ConVulPOE, a dynamic analysis targeting memory bugs due to concurrency [42]. Similarly to M2 and SeqCheck, ConVulPOE observes bug-free traces, and attempts to predict valid reorderings that will expose a memory bug, such as null pointer dereferences. In Figure 14, trace σ is bug-free, while σ' is a valid permutation of σ that exposes a null-pointer dereference. Internally, it utilizes STs for handling reachability queries. The reachability queries are incremental. We measured the total time taken in the analysis.

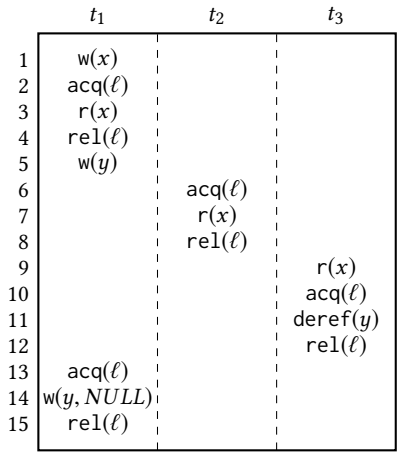

(a) Trace σ doesn't have a null-pointer dereference.

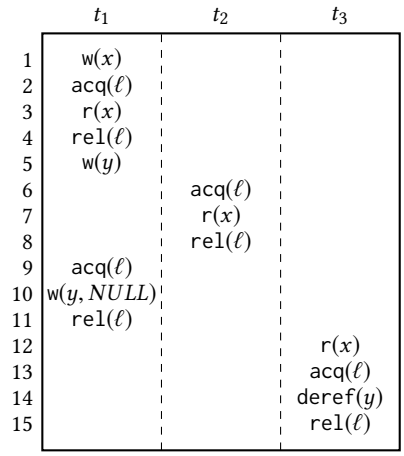

(b) In trace σ', (e_{10}, e_{14}) forms a null-pointer dereference.

Fig. 14. An example of memory bugs, taken from [42]. On the left, trace σ is a valid trace, while on the right, trace σ' is a permutation of σ with a memory bug. e_{14} dereferencing y written by e_{10} leads to a null-pointer dereference. Similarly to Figures 12 and 13, the analysis guarantees that σ' is a valid trace because every read event observes the same write event as in σ.

Table 4. Memory Bug Prediction Results

benchmark	T	N	q	VCs (s)	STs† (s)	CSSTs (s)
pbzip2	7	243.5K	0.06	298	6.7	3.8
pigz	6	2.3M	0.01	T.O.	2087	1265
xz	2	3.0M	0.01	15.1	13.9	12.9
lbzip2	11	9.3M	0.04	273	96.6	46.9
x264	7	10.4M	0.03	474	165	62.0
libvpx	2	19.0M	0.01	440	83.6	77.6
libwebp	2	29.5M	0.1	324	118	115
x265	15	37.1M	0.02	T.O.	O.O.M.	627
Total	**-**	**110.9M**	**-**	**>37824**	**>20570**	**2211**

The timeout is 5h. Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

Our results are shown in Table 4. We again observe that VCs are the slowest data structure to support this analysis. While STs offer a generous speedup, CSSTs are again the most performant, and even manage to handle benchmarks that makes VCs to time out and STs to go out of memory.

4. Consistency checking in x86-TSO. Here we report our results for consistency checking under the x86-TSO memory model. In Figure 15, the trace σ is an incorrect execution for x86-TSO. Though the problem is NP-complete, polynomial-time heuristics are known to be remarkably accurate. We experiment with the consistency analysis developed in [35]. In contrast to the previous analyses where the underlying chain DAG had one chain per thread, here we have two chains per thread, one for its program order and one for its store buffer, as per the operational semantics of x86-TSO [37]. The reachability queries are incremental. We measure the total time taken in the analysis.

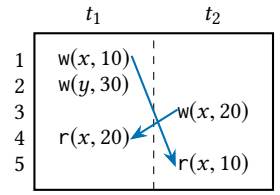

Fig. 15. The trace σ forms an inconsistent x86 execution, taken from [35]. In particular, the rf edge from e_1 to e_5 is illegal according to TSO semantics. Instead, e_5 can only read from the write event e_3. The analysis infers the potential writes that a read may read from according to the x86 semantics by applying various partial order-based rules.

Table 5. x86-TSO Consistency Checking Results

benchmark	T	N	q	VCs (s)	STs (s)	CSSTs (s)
dekker	3	16.3K	0.41	6.9	0.3	0.25
peterson	3	19.0K	0.37	9.4	0.4	0.45
lamport	3	28.1K	0.39	18.3	0.6	0.6
dq	4	31.5K	0.24	11.5	0.5	0.44
chase-lev	5	32.2K	0.24	23.3	0.6	0.51
szymanski	3	77.9K	0.33	17.4	0.9	0.8
buf-ring	9	115.3K	0.39	322	7.3	6.27
mcs-lock	11	196.4K	0.17	298	28.3	20.71
spsc	3	243.6K	0.53	2548	2.1	1.74
linuxrwlocks	6	276.4K	0.28	T.O.	12.6	11.18
fib-bench	3	300.0K	0.4	T.O.	4.7	4.21
seqlock	17	318.1K	0.15	1227	46.1	28.16
spinlock	11	482.5K	0.39	T.O.	103	75.11
ttaslock	11	491.1K	0.41	T.O.	111	81.19
exp-bug	4	498.5K	0.38	2591	3.6	2.76
mutex	11	519.7K	0.56	T.O.	122	86.09
ticketlock	6	569.6K	0.4	T.O.	20.8	17.17
gcd	3	750.1K	0.28	T.O.	6.0	4.52
indexer	17	800.0K	0.51	7.8	5.8	3.23
twalock	11	900.0K	0.43	T.O.	174	118.39
treiber	6	1.0M	0.22	T.O.	20.7	16.73
mpmc	10	2.0M	0.17	T.O.	41.6	22.59
barrier	5	2.8M	0.6	T.O.	139	112.61
Total	-	**12.5M**	-	**>46680**	**852**	**616**

The timeout is 1h. Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

The results are shown in Table 5. VCs are very slow compared to the other two data structures. Indeed, this analysis is quite demanding, as it performs repeated updates between events that are in the middle of the partial order, which leads to deep edge propagations. Although STs perform considerably better than VCs, CSSTs are decisively faster in this setting too. Finally, the average density q is larger here than in previous analyses, but still considerably below 1.

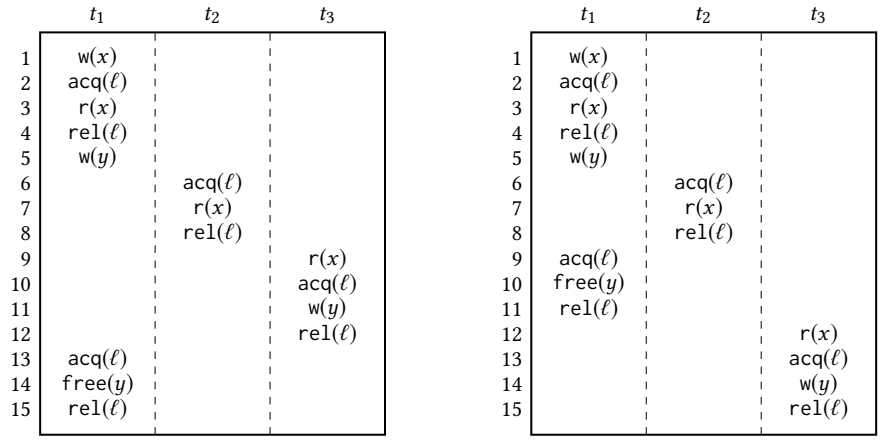

(a) Trace σ doesn't have a use-after-free bug. (b) In trace σ', (e_{10}, e_{14}) forms a use-after-free.

Fig. 16. An example of a use-after-free bug that occurs in a concurrent context. On the left, trace σ is a valid trace, while on the right, trace σ' is a permutation of σ with a use-after-free bug. This is formed due to the event pair σ', (e_{10}, e_{14}), where the former frees the location y and the latter attempts to write to the freed location.

Table 6. Use-after-free Results

benchmark	T	N	q	VCs† (s)	STs (s)	CSSTs (s)
bbuf	3	27.7K	0.02	426	459	415
BoundedBuffer	11	325.7K	0.02	835	736	679
DiningPhil	21	1.4M	0.01	1559	1191	1110
fanger01-ok	5	93.4K	0.06	366	194	190
qtsort	6	41.7M	0.01	104	48.2	<u>31.9</u>
pbzip	5	269.3K	0.01	T.O.	15894	14471
Total	-	**43.8M**	-	**>21290**	**18521**	**16898**

The timeout is 5h. Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

5. Use-after-free prediction. Here we incorporate CSSTs in the UFO analysis for use-after-free vulnerabilities due to concurrency [20] Figure 16 shows such an example. This analysis is based on SMT, but relies on efficient partial-order reasoning to generate SMT queries. The reachability queries are incremental. We measure the time to generate the queries. Our results are shown in Table 6.

In alignment with our observations so far, CSSTs outperform VCs and STs on all benchmarks. However the speedup is not as large as in other analyses. Looking closer into the running times, we observe that the analysis itself spends considerable time in components other than maintaining reachability, which are common for all data structures. If we only focus on the time taken for maintaining the partial order, CSSTs indeed offer a significant speedup, e.g., 3.5× and 1.7× in BoundedBuffer over VCs and STs, respectively.

6. Data-race detection in C11. Here we incorporate CSSTs in C11Tester, a data-race detector for the C11 memory model [24]. The analysis constructs incrementally a trace σ, by iteratively mapping a read to a write to obtain its value from. Figure 17 shows an example of how C11Tester picks up different values for the same read event. To make these choices consistent with the C11

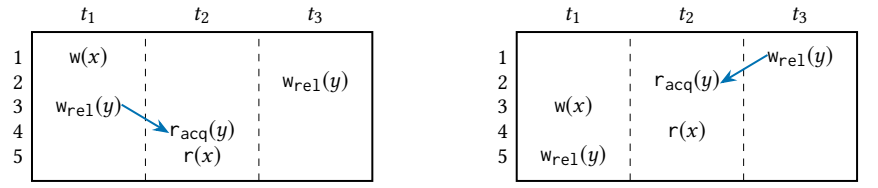

(a) In trace σ, (e_1, e_5) is not a race. (b) In trace σ', (e_3, e_4) forms a race.

Fig. 17. C11Tester can reschedule a race-free trace to explore potential races. On the left, σ is race free because (e_1, e_5) are ordered by the rf edge from e_3 to e_4. On the right, σ' is a permutation of σ that exposes a data race.

memory model, C11Tester maintains a partial order P using VCs. In high level, P consists of several, independent partial orders $\{P_x\}_x$, one for each memory location x that is accessed by the program. Upon pairing a read event with a write event, a set of other write events on the same location are ordered before the writer of the read, and these orderings are further propagated to the successors of these write events. The reachability queries are incremental. To ensure that all data structures process the same trace, we run them alongside each other on a single execution, and measure the time taken to maintain the partial order in each.

Our results are shown in Table 7. Interestingly, VCs are generally more performant than STs and CSSTs in this setting. Looking closer into each benchmark, we observe that new orderings lead to very little (typically none at all) propagation in P. Effectively, this makes VCs run in $O(1)$ time per update, and thus gain an advantage. In contrast, in readerswriters and atomicblocks C11Tester inserts non-trivial orderings in P (i.e., orderings that lead to propagation), and CSSTs are considerably more performant than VCs. Finally, the average density q is high in a few benchmarks, reaching even the value 1, which explains why CSSTs and STs have comparable performance, though CSSTs remain more performant.

7. View serializability. Here we incorporate CSSTs in a view serializability [6] checking engine. This analysis observes a trace σ that has certain sections marked as atomic. Atomic sections can overlap in general, and the task of the analysis is to prove σ serializable, meaning that it can be permuted to an equivalent, serial trace in which atomic sections do not overlap. Figure 18 illustrates an example of a serializable trace and its serial-equivalent trace. To discover this reordering, the analysis iteratively orders events according to several saturation rules. The reachability queries are incremental, and we measure the total time required for the entire analysis.

Our results are shown in Table 8. We observe that CSSTs are the fastest data structure to support this analysis, while the performance gap between STs and VCs is not significant. In contrast to previous settings, here we also observe that for some benchmarks, such as series and batik, STs run slower than VCs. In these cases, we observed that STs spend most of its time in initializing the underlying ST data structure. Notably, the average density q is very low in these benchmarks. This is because the analysis determines that the trace is not view serializable early, without populating CSSTs with many edges. This results in CSSTs significantly outperforming VCs and STs in this analysis.

8. Root-causing linearizability violations. Here we incorporate CSSTs in a recent root-cause analysis for linearizability violations [13]. Figure 19 displays a linearizability violation, and the goal of this analysis is to identify the root causes of such violations. The analysis achieves this by finding optimal repairs that rule out non-linearizable executions while minimizing the exclusion of linearizable ones. The reachability queries are both incremental and decremental, hence the only baseline for maintaining the partial order is plain Graphs, as indeed used in [13]. We measure the total time taken in the analysis.

Table 7. Race Detection on C11

benchmark	T	N	q	VCs† (s)	STs (s)	CSSTs (s)
dq	5	72.8K	0.48	2.5	2.8	3.1
mabain	7	101.6K	0.33	0.6	1.2	1.1
seqlock	18	693.9K	0.2	4.5	6.7	6.3
iris-1	13	1.1M	0.2	54.8	66.5	60.3
qu	11	1.3M	0.43	5.9	7.5	7.2
indexer	18	1.6M	1	1.0	1.1	1.1
exp-bug	5	2.5M	0.25	2.6	2.5	2.5
twalock	12	4.5M	0.34	29.6	53.3	49.0
gcd	4	4.5M	1	2.6	2.8	2.7
spinlock	12	4.8M	0.2	23.0	37.9	34.0
ttaslock	12	4.9M	0.25	26.7	48.0	42.7
silo	5	5.6M	0.01	5.9	6.2	6.1
fib-bench	4	6.0M	1	6.5	6.3	6.5
linuxrwlocks	7	6.9M	0.39	30.8	40.6	37.7
barrier	6	8.3M	0.51	19.6	24.5	23.7
mpmc	11	9.2M	0.24	37.6	44.9	42.6
spsc	4	9.7M	0.55	12.2	12.2	12.1
mcs-lock	12	9.9M	0.37	55.2	86.8	80.1
treiber	7	10.1M	0.11	31.3	36.6	35.7
iris-2	4	11.5M	0.2	15.5	14.7	14.7
gdax	8	13.5M	0.97	8.5	7.6	7.6
ticketlock	7	14.2M	0.48	45.1	70.4	64.2
mutex	12	15.6M	0.39	48.5	68.5	63.7
readerswriters	13	10.1M	0.33	16.3	7.6	7.5
atomicblocks	33	15.5M	0.5	9.7	1.7	1.6
Total	-	**172.4M**	-	**496**	**659**	**614**

Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

Our results are shown in Table 9. The experimental setup of [13] consists of three data structures, accessed an increasing number of times. CSSTs are typically orders of magnitude faster than plain Graphs used in [13], which strongly supports CSSTs as an efficient data structure also for analyses utilizing decremental partial-order updates.

5.3 Scalability Analysis

Here we perform controlled scalability experiments on the performance of CSSTs for edge insertions and queries. We consider partial orders P in two settings, consisting of $k = 10, 20$ chains, and varying number of events ℓ in each chain. Initially, P has no cross-chain edges. First, to measure the performance of edge insertions, we repeatedly insert random cross-chain edges $\langle t, i\rangle \mapsto \langle t', j\rangle$ such that the two endpoints are unordered and $|i - j| \leq b$. The last condition captures the fact that cross-chain orderings are typically between events that execute within the same time-window. We report on window $b = 10^4$, though we have observed similar results with other values of b. This condition also prevents P from becoming close to total. We attempt to insert 20ℓ edges in each case, and report the average time of edge insertions. Second, to measure the performance of edge

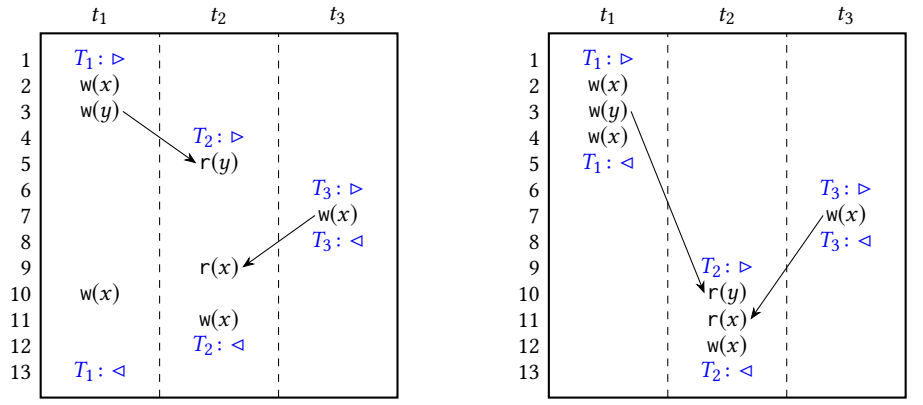

(a) σ is view serializable. (b) Serial trace σ' is a permutation of σ.

Fig. 18. An example of view serializability analysis. View-serializability ensures that a program produces the same final results as an equivalent serial execution, even if the transactions interleave in the original trace. Trace σ is view serializable, which means that it can be rescheduled to a serial trace contains the same rf relations. On the right, σ' is such a serial trace that is a valid permutation of σ.

Table 8. View Serializability Checking Results

benchmark	T	N	q	VCs (s)	STs (s)	CSSTs (s)
philo	6	4.4k	0.08	0.03	0.03	0.03
hedc	7	86k	<0.01	0.08	0.06	0.05
lufact	4	16M	<0.01	0.02	0.01	0.01
pmd	13	20M	<0.01	1.65	1.55	0.691
tsp	9	22M	0.02	1.52	1.45	0.02
sor	4	39M	<0.01	3.18	2.83	2.27
tomcat	48	59M	<0.01	6.32	3.70	0.05
sunflow	16	128M	<0.01	8.99	9.03	0.02
xalan	13	158M	0.02	3.17	2.23	1.15
series	4	306M	<0.01	2.89	12.50	0.004
batik	6	397M	<0.01	13.07	14.07	3.25
crypt	7	458M	<0.01	T.O.	O.O.M.	41.42
Total	**-**	**5.87G**	**-**	**>7241**	**>7248**	**3649**

The timeout is 1h. Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

queries, we use the partial order P at the end of the above process. We make 10^6 random queries, and report the average time per query.

The results are shown in Figure 20. As expected by theory, VCs suffer linear complexity for insertions but achieve queries in constant time. On the other hand, CSSTs and STs achieve logarithmic complexity in both insertions and deletions. CSSTs are far more performant than VCs in edge insertions, but also achieve query times that are similar to VCs, though, naturally, VCs are slightly faster, as each query is a simple lookup. This is achieved by our minima indexing and sparse representation, which makes queries run with minimal overhead in practice. Finally, CSSTs are

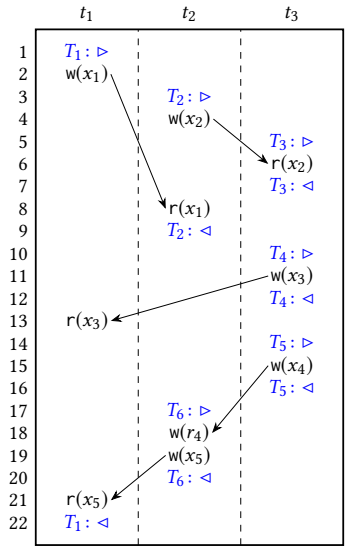

Fig. 19. The trace σ witnesses a linearizability violation. The root cause analysis computes repairs and performs a prioritization among them so as to minimize the amount of synchronization introduced. Here, while the violation can be fixed by introducing atomic sections either between ($\mathtt{w}(x_1)$, $\mathtt{r}(x_3)$) or ($\mathtt{w}(x_1)$, $\mathtt{r}(x_5)$), the former repair is preferred because it is subsumed by the other larger atomic section that includes $\mathtt{r}(x_5)$.

Table 9. Root Causing Linearizability Violation Results

benchmark	T	N	q	Graphs† (s)	CSSTs (s)
LogicalOrderingAVL	3	615	0.1	1.2	0.4
	3	1.6K	0.07	6.1	1.6
	3	2.6K	0.06	13.5	3.8
	3	5.0K	0.05	49.5	10.8
OptimisticListSorted-SetWaitFreeContains	3	219	0.19	1.4	0.2
	3	399	0.18	9.8	0.4
	3	759	0.17	114	0.9
	3	1.1K	0.17	512	1.7
RWLockCoarse-GrainedListIntSet	3	978	0.04	3.6	0.8
	3	1.6K	0.03	11.4	1.7
	3	3.2K	0.03	73.6	5.3
	3	4.8K	0.03	249	10.1
Total	-	**22.9K**	-	**1045**	**38**

Underlined values represent the best-performing data structure by total runtime, with at least a 10% improvement over the compared methods.

considerably faster than STs in both insertions and queries, an improvement stemming from the SST (Section 3.2).

6 Related Work

The efficient maintenance of partial orders is a key component in concurrent program analysis. Vector Clocks [29] were originally proposed as an efficient mechanism for causality tracking

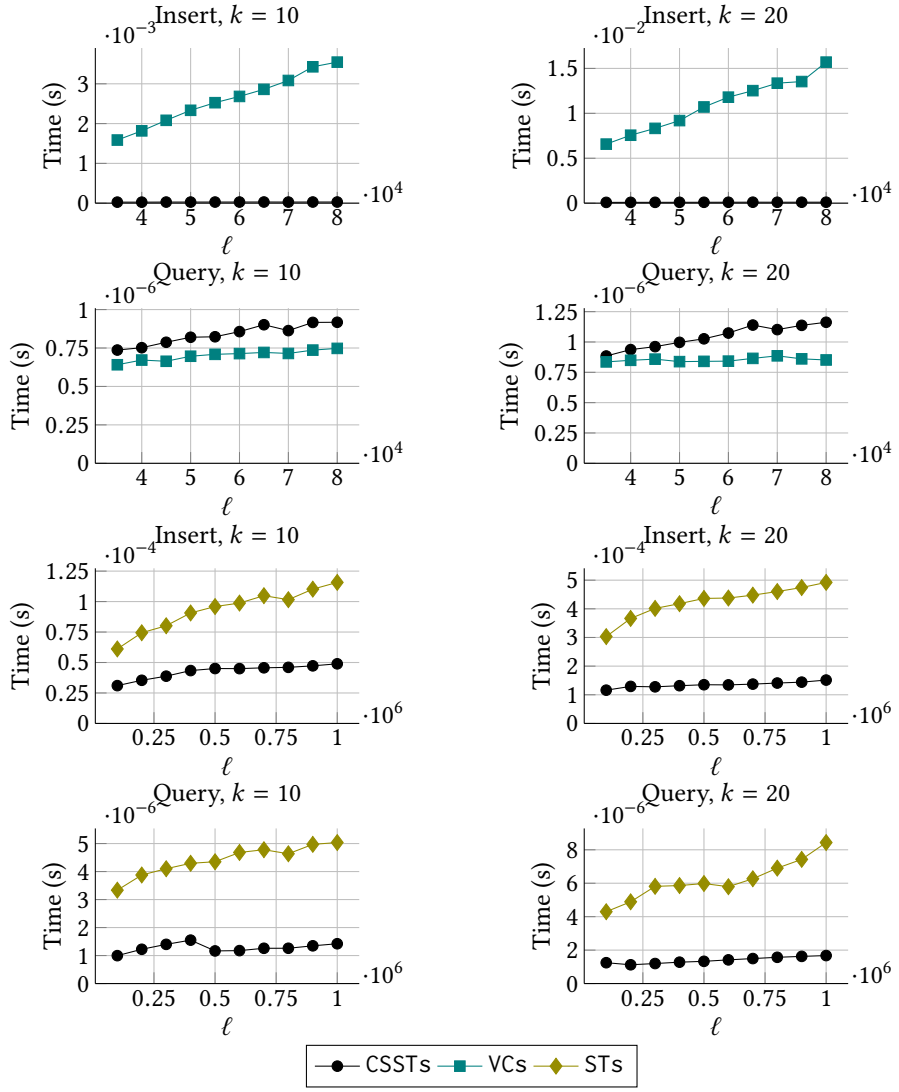

Fig. 20. Scalability experiments on CSSTs, STs , and VCs. Each partial order contains k chains and $\ell = n/k$ events per chain.

in distributed systems [22], and have been used extensively in program analysis. Their efficient implementation has been approached in various ways, e.g., using varying-size arrays [12], AVL trees [23], Chain Clocks [3], and Tree Clocks [26]. These variations offer performance improvements over vanilla Vector Clocks, but only when the partial order is built incrementally and in a streaming fashion.

For non-streaming updates, partial orders are normally represented as plain graphs [2, 7, 13, 31, 34, 43] and Vector Clocks [18, 24]. The closest data structure to CSSTs has been STs, developed in [32] for data-race detection, and subsequently used in other dynamic analyses [9, 38, 42]. The development of CSSTs lies on a few technical novelties. For example, our minima indexing and sparse tree representation (Section 3.2) are novel, and lead to tighter complexity bounds. Moreover,

in order to handle the fully dynamic setting, CSSTs do not store transitive reachability. This requires overcoming the challenge of discovering transitive reachability during queries in an efficient manner (Section 3.3).

Besides the eight applications we have experimented with in Section 5, CSSTs can also be directly utilized in other types of analyses. One such example is in the saturation procedure for consistency checking executions against PSO, for driving stateless model checking [8]. Here, each chain in the CSSTs would capture either the program order of a thread, or the store buffer of a thread for a specific memory location. Another example is testing database histories for strong isolation levels such as prefix consistency and snapshot isolation [6], where each chain in the CSSTs captures the session order of a set of transactions. A third example is the vindication phase of the Vindicator data race predictor [34].

7 Conclusion

We have introduced CSSTs, a data structure for the fully-dynamic maintenance of partial orders with small width, supporting a plethora of dynamic analyses for concurrency. Our experimental results indicate that CSSTs are an efficient, drop-in replacement of existing approaches to maintaining partial orders in various application domains of recent literature. Our aim is for CSSTs to become a standard reference for driving future scalable dynamic analyses.

There are also several interesting future research directions. First, it would be beneficial to make CSSTs concurrent/thread-safe, so that they can be used in an online monitoring setting, such as TSAN [36]. Second, lower-level implementation details may further reduce the running time of CSSTs. For example, instead of storing outgoing edges from thread t to t', each CSST may store incoming edges to t' from t. The overall representation of CSSTs remains the same, though the suffix-minima queries must be replaced by prefix-maxima queries. This may reduce the cross-chain density of each thread, and thus reduce further the time required to perform CSST-related operations (insert/delete/query).

References

[1] 2024. CSSTs GitHub Repository. https://github.com/hcantunc/cssts. Last accessed: November 16, 2025.

[2] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, Magnus Lång, Tuan Phong Ngo, and Konstantinos Sagonas. 2019. Optimal stateless model checking for reads-from equivalence under sequential consistency. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 29. DOI: https://doi.org/10.1145/3360576

[3] Anurag Agarwal and Vijay K. Garg. 2005. Efficient dependency tracking for relevant events in shared-memory systems. In *Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC'05)*. ACM, New York, NY, USA, 19–28. DOI: https://doi.org/10.1145/1073814.1073818

[4] Pratyush Agarwal, Krishnendu Chatterjee, Shreya Pathak, Andreas Pavlogiannis, and Viktor Toman. 2021. Stateless model checking under a reads-value-from equivalence. In *Computer Aided Verification*, Alexandra Silva and K. Rustan M. Leino (Eds.). Vol. 12759, Springer International Publishing, Cham, 341–366. DOI: https://doi.org/10.1007/978-3-030-81685-8_16

[5] Lars Arge, Johannes Fischer, Peter Sanders, and Nodari Sitchinava. 2013. On (dynamic) range minimum queries in external memory. In *Algorithms and Data Structures*, Frank Dehne, Roberto Solis-Oba, and Jörg-Rüdiger Sack (Eds.). Springer Berlin Heidelberg, Berlin, 37–48. DOI: https://doi.org/10.1007/978-3-642-40104-6_4

[6] Ranadeep Biswas and Constantin Enea. 2019. On the complexity of checking transactional consistency. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 165:1–165:28. DOI: https://doi.org/10.1145/3360591

[7] Swarnendu Biswas, Jipeng Huang, Aritra Sengupta, and Michael D. Bond. 2014. DoubleChecker: Efficient sound and precise atomicity checking. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'14)*. ACM, New York, NY, USA, 28–39. DOI: https://doi.org/10.1145/2594291.2594323

[8] Truc Lam Bui, Krishnendu Chatterjee, Tushar Gautam, Andreas Pavlogiannis, and Viktor Toman. 2021. The reads-from equivalence for the TSO and PSO memory models. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–30. DOI: https://doi.org/10.1145/3485541

[9] Yan Cai, Hao Yun, Jinqiu Wang, Lei Qiao, and Jens Palsberg. 2021. Sound and efficient concurrency bug prediction. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the*

Foundations of Software Engineering (ESEC/FSE 2021). ACM, New York, NY, USA, 255–267. DOI: https://doi.org/10.1145/3468264.3468549

[10] Soham Chakraborty, Shankara Narayanan Krishna, Umang Mathur, and Andreas Pavlogiannis. 2024. How hard is weak-memory testing? *Proc. ACM Program. Lang.* 8, POPL (2024), 66:1978–66:2009. DOI: https://doi.org/10.1145/3632908

[11] Marek Chalupa, Krishnendu Chatterjee, Andreas Pavlogiannis, Nishant Sinha, and Kapil Vaidya. 2018. Data-centric dynamic partial order reduction. *Proc. ACM Program. Lang.* 2, POPL (2018), 1–30. DOI: https://doi.org/10.1145/3158119

[12] Mark Christiaens and Koenraad De Bosschere. 2001. Accordion clocks: Logical clocks for data race detection. In *Proceedings of the 7th International Euro-Par Conference Manchester on Parallel Processing (Euro-Par'01)*. Springer-Verlag, Berlin, 494–503. DOI: https://doi.org/10.1007/3-540-44681-8_73

[13] Berk Çirisci, Constantin Enea, Azadeh Farzan, and Suha Orhun Mutluergil. 2020. Root causing linearizability violations. In *Computer Aided Verification*. Shuvendu K. Lahiri and Chao Wang (Eds.), Springer International Publishing, Cham, 350–375. DOI: https://doi.org/10.1007/978-3-030-53288-8_17

[14] Ariel Eizenberg, Yuanfeng Peng, Toma Pigli, William Mansky, and Joseph Devietti. 2017. BARRACUDA: Binary-level analysis of runtime races in cuda programs. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2017)*. ACM, New York, NY, USA, 126–140. DOI: https://doi.org/10.1145/3062341.3062342

[15] Michael Emmi and Constantin Enea. 2019. Weak-consistency specification via visibility relaxation. *Proc. ACM Program. Lang.* 3, POPL (2019), 28. DOI: https://doi.org/10.1145/3290373

[16] Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: Efficient and precise dynamic race detection. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'09)*. ACM, New York, NY, USA, 121–133. DOI: https://doi.org/10.1145/1542476.1542490

[17] Cormac Flanagan, Stephen N. Freund, and Jaeheon Yi. 2008. Velodrome: A sound and complete dynamic atomicity checker for multithreaded programs. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'08)*. ACM, New York, NY, USA, 293–303. DOI: https://doi.org/10.1145/1375581.1375618

[18] Mingyu Gao, Soham Chakraborty, and Burcu Kulahcioglu Ozkan. 2023. Probabilistic concurrency testing for weak memory programs. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS 2023)*. ACM, New York, NY, USA, 603–616. DOI: https://doi.org/10.1145/3575693.3575729

[19] Phillip B. Gibbons and Ephraim Korach. 1997. Testing Shared Memories. *SIAM J. Comput.* 26, 4, 1208–1244. DOI: https://doi.org/10.1137/S0097539794279614

[20] Jeff Huang. 2018. UFO: Predictive concurrency use-after-free detection. In *Proceedings of the 40th International Conference on Software Engineering (ICSE'18)*. ACM, New York, NY, USA, 609–619. DOI: https://doi.org/10.1145/3180155.3180225

[21] Christian Gram Kalhauge and Jens Palsberg. 2018. Sound deadlock prediction. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 29. DOI: https://doi.org/10.1145/3276516

[22] Leslie Lamport. 1978. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* 21, 7 (1978), 558–565. DOI: https://doi.org/10.1145/359545.359563

[23] Guangpu Li, Shan Lu, Madanlal Musuvathi, Suman Nath, and Rohan Padhye. 2019. Efficient scalable thread-safety-violation detection: Finding thousands of concurrency bugs during testing. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP'19)*. ACM, New York, NY, USA, 162–180. DOI: https://doi.org/10.1145/3341301.3359638

[24] Weiyu Luo and Brian Demsky. 2021. C11Tester: A race detector for C/C++ atomics. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'21)*. ACM, New York, NY, USA, 630–646. https://doi.org/10.1145/3445814.3446711

[25] Umang Mathur, Dileep Kini, and Mahesh Viswanathan. 2018. What happens-after the first race? Enhancing the predictive power of happens-before based dynamic race detection. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 29. DOI: https://doi.org/10.1145/3276515

[26] Umang Mathur, Andreas Pavlogiannis, Hünkar Can Tunç, and Mahesh Viswanathan. 2022. A tree clock data structure for causal orderings in concurrent executions. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'22)*. ACM, New York, NY, USA, 710–725. DOI: https://doi.org/10.1145/3503222.3507734

[27] Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. 2020. The complexity of dynamic data race prediction. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*. ACM, 713–727. DOI: https://doi.org/10.1145/3373718.3394783

[28] Umang Mathur and Mahesh Viswanathan. 2020. Atomicity checking in linear time using vector clocks. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'20)*. ACM, New York, NY, USA, 183–199. DOI: https://doi.org/10.1145/3373376.3378475

[29] Friedemann Mattern. 1989. Virtual time and global states of distributed systems. In *Parallel and Distributed Algorithms: Proceedings of the International Workshop on Parallel & Distributed Algorithms*, M. Cosnard et al. (Ed.). Elsevier Science Publishers B. V., 215–226.

[30] Madanlal Musuvathi, Shaz Qadeer, Thomas Ball, Gerard Basler, Piramanayagam Arumuga Nainar, and Iulian Neamtiu. 2008. Finding and reproducing heisenbugs in concurrent programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation (OSDI'08)*. USENIX Association, USA, 267–280.

[31] Brian Norris and Brian Demsky. 2013. CDSchecker: Checking concurrent data structures written with C/C++ atomics. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications (OOPSLA'13)*. ACM, New York, NY, USA, 131–150. DOI: https://doi.org/10.1145/2509136.2509514

[32] Andreas Pavlogiannis. 2019. Fast, sound, and effectively complete dynamic race prediction. *Proc. ACM Program. Lang.* 4, POPL (2019), 29. DOI: https://doi.org/10.1145/3371085

[33] Jake Roemer, Kaan Genç, and Michael D. Bond. 2020. SmartTrack: Efficient predictive race detection. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. ACM, New York, NY, USA, 747–762. DOI: https://doi.org/10.1145/3385412.3385993

[34] Jake Roemer, Kaan Genç, and Michael D. Bond. 2018. High-coverage, unbounded sound predictive race detection. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2018)*. ACM, New York, NY, USA, 374–389. DOI: https://doi.org/10.1145/3192366.3192385

[35] Amitabha Roy, Stephan Zeisset, Charles J. Fleckenstein, and John C. Huang. 2006. Fast and generalized polynomial time memory consistency verification. In *Computer Aided Verification*, Thomas Ball and Robert B. Jones (Eds.). Springer Berlin Heidelberg, Berlin, 503–516. DOI: https://doi.org/10.1007/11817963_46

[36] Konstantin Serebryany and Timur Iskhodzhanov. 2009. ThreadSanitizer: Data race detection in practice. In *Proceedings of the Workshop on Binary Instrumentation and Applications (WBIA'09)*. ACM, New York, NY, USA, 62–71. DOI: https://doi.org/10.1145/1791194.1791203

[37] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. 2010. X86-TSO: A rigorous and usable programmer's model for X86 multiprocessors. *Commun. ACM* 53, 7 (2010), 89–97. DOI: https://doi.org/10.1145/1785414.1785443

[38] Zheng Shi, Umang Mathur, and Andreas Pavlogiannis. 2024. Optimistic prediction of synchronization-reversal data races. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE'24)*. Association for Computing Machinery, New York, NY, USA, Article 134, 13 pages. https://doi.org/10.1145/3597503.3639099

[39] Hünkar Can Tunç, Parosh Aziz Abdulla, Soham Chakraborty, Shankaranarayanan Krishna, Umang Mathur, and Andreas Pavlogiannis. 2023. Optimal reads-from consistency checking for C11-style memory models. *Proc. ACM Program. Lang.* 7, PLDI (2023), 137:761–137:785. DOI: https://doi.org/10.1145/3591251

[40] Hünkar Can Tunç, Ameya Prashant Deshmukh, Berk Cirisci, Constantin Enea, and Andreas Pavlogiannis. 2024. CSSTs: A dynamic data structure for partial orders in concurrent execution analysis. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'24, Vol. 3)*. ACM, New York, NY, USA, 223–238. DOI: https://doi.org/10.1145/3620666.3651358

[41] Hünkar Can Tunç, Umang Mathur, Andreas Pavlogiannis, and Mahesh Viswanathan. 2023. Sound dynamic deadlock prediction in linear time. *Proc. ACM Program. Lang.* 7, PLDI (2023), 26. DOI: https://doi.org/10.1145/3591291

[42] Kunpeng Yu, Chenxu Wang, Yan Cai, Xiapu Luo, and Zijiang Yang. 2021. Detecting concurrency vulnerabilities based on partial orders of memory and thread events. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2021)*. ACM, New York, NY, USA, 280–291. DOI: https://doi.org/10.1145/3468264.3468572

[43] Rachid Zennou, Mohamed Faouzi Atig, Ranadeep Biswas, Ahmed Bouajjani, Constantin Enea, and Mohammed Erradi. 2020. Boosting sequential consistency checking using saturation. In *Automated Technology for Verification and Analysis*, Dang Van Hung and Oleg Sokolsky (Eds.). Springer International Publishing, Cham, 360–376. DOI: https://doi.org/10.1007/978-3-030-59152-6_20

[44] Rachid Zennou, Ahmed Bouajjani, Constantin Enea, and Mohammed Erradi. 2019. Gradual consistency checking. In *Computer Aided Verification*, Isil Dillig and Serdar Tasiran (Eds.). Springer International Publishing, Cham, 267–285. DOI: https://doi.org/10.1007/978-3-030-25543-5_16

Received 7 February 2025; revised 26 July 2025; accepted 18 September 2025