

PaC-trees: Supporting Parallel and Compressed Purely-Functional Collections

Laxman Dhulipala
University of Maryland
laxman@umd.edu

Guy Blelloch
Carnegie Mellon University
guyb@cs.cmu.edu

Yan Gu
UC Riverside
ygu@cs.ucr.edu

Yihan Sun
UC Riverside
yihans@cs.ucr.edu

Abstract

Many modern programming languages are shifting toward a functional style for collection interfaces such as sets, maps, and sequences. Functional interfaces offer many advantages, including being safe for parallelism and providing simple and lightweight snapshots. However, existing high-performance functional interfaces such as PAM, which are based on balanced purely-functional trees, incur large space overheads for large-scale data analysis due to storing every element in a separate node in a tree.

This paper presents PaC-trees, a purely-functional data structure supporting functional interfaces for sets, maps, and sequences that provides a significant reduction in space over existing approaches. A PaC-tree is a balanced binary search tree which blocks the leaves and compresses the blocks using arrays. We provide novel techniques for compressing and uncompressing the blocks which yield practical parallel functional algorithms for a broad set of operations on PaC-trees such as union, intersection, filter, reduction, and range queries which are both theoretically and practically efficient.

Using PaC-trees we designed CPAM, a C++ library that implements the full functionality of PAM, while offering significant extra functionality for compression. CPAM consistently matches or outperforms PAM on a set of microbenchmarks on sets, maps, and sequences while using about a quarter of the space. On applications including inverted indices, 2D range queries, and 1D interval queries, CPAM is competitive with or faster than PAM, while using 2.1–7.8x less space. For static and streaming graph processing, CPAM offers 1.6x faster batch updates while using 1.3–2.6x less space than the state-of-the-art graph processing system Aspen.

1 Introduction

Almost all modern programming languages include extensive support for collections, such as sets, maps, and sequences either as libraries or built-in data types. Support for such collections has become the cornerstone of large-scale data processing, as exemplified by systems such as Apache Spark [54]. Among the interfaces for collections, there has been a trend towards a functional style, shying away from mutation (e.g.,

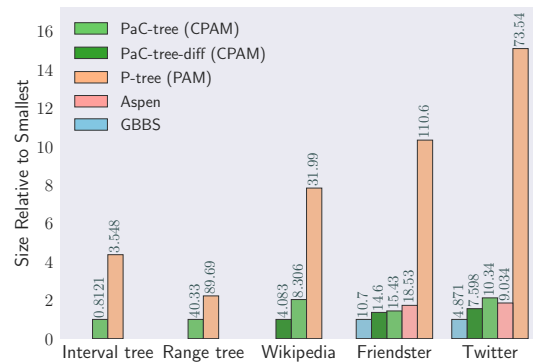


Figure 1. Relative sizes of the interval tree, range tree, inverted index (Wikipedia corpus), and graph representations (Twitter, Friendster) studied in this paper using PaC-trees from CPAM (using $B = 128$) and other systems. Lower is better. The numbers shown on top of the bars are the sizes of each representation in GiB. PaC-tree-diff compresses integer keys using difference encoding. The C-trees from Aspen [25] also support difference encoding. GBBS is the *static* compressed graph representation from the Graph Based Benchmark Suite [26] which uses difference encoding, and serves as a baseline for the tree-based graph representations.

Spark is functional). Functional interfaces have several advantages over mutating ones, including being safe for parallelism, allowing safe composition, permitting flexible implementations (e.g., using copies when helpful), and supporting snapshots. Supporting snapshots is particularly useful in scenarios in which a stream of updates is being made to a collection which is concurrently being analyzed [22, 25, 37, 42].

Recent work [52] has developed a purely functional library, PAM, for representing sequences, ordered sets, ordered maps, and augmented maps (defined in [52]) using balanced trees, called *P-trees*. P-trees use path copying to perform updates, supporting functional updates at a reasonably low cost (e.g., $O(\log n)$ per point update). However they come at a cost of high space usage—every element requires a node in the tree. This is particularly problematic for large-scale data analysis, since in large-systems memory is often the dominating cost.

In this paper we present **Parallel Compressed trees (PaC-trees)**: a purely-functional data structure for supporting a similar functionality as P-trees but with significant reduction in space—up to an order of magnitude (see Fig. 1). Our approach is based on blocking the leaves and compressing the blocks using arrays (see Fig. 4). We present innovative techniques for compressing and uncompressing the blocks

	Primitive	Work	Span
Sequence	<i>Build</i>	$O(n)$	$O(\log n)$
	<i>Map</i>	$O(n)$	$O(\log n)$
	<i>Filter</i>	$O(n)$	$O(\log n)$
	<i>Reduce</i>	$O(n)$	$O(\log n)$
	<i>Take</i>	$O(\log n + B)$	$O(\log n)$
	<i>n-th</i>	$O(\log n + B)$	$O(\log n)$
	<i>FindFirst</i>	$O(k)$	$O(\log n)$
	<i>Append[†]</i>	$O(\log n + B)$	$O(\log n)$
	<i>Reverse[†]</i>	$O(n)$	$O(\log n)$
Set and Map	<i>Build</i>	$O(n \log n)$	$O(\log n)$
	<i>Next/Previous</i>	$O(\log n + B)$	$O(\log n)$
	<i>Rank</i>	$O(\log n + B)$	$O(\log n)$
	<i>Range</i>	$O(\log n + B)$	$O(\log n)$
	<i>Insert</i>	$O(\log n + B)$	$O(\log n)$
	<i>Union</i>	$O(m \log \frac{n}{m} + \min(mB, n))$	$O(\log n \log m)$
	<i>Intersect</i>	$O(m \log \frac{n}{m} + \min(mB, n))$	$O(\log n \log m)$
	<i>Difference</i>	$O(m \log \frac{n}{m} + \min(mB, n))$	$O(\log n \log m)$

Table 1. Primitives from the Sequence, Set, and Map interfaces in CPAM, including the work and span bounds. Note that primitives marked with [†] are specific to Sequences, and Set and Map primitives cannot be applied to Sequences. m, n are defined to be the size of the smaller and larger sets, respectively. B is the block size (the size of a blocked leaf in a PaC-tree). We assume a parallelizable encoding for the span bounds.

without needing to re-implement the full functionality of P-trees. Importantly, in the paper we analyze the cost of all the operations as a function of the block size B as well as the collection size. This is analyzed both in terms of the work (runtime sequentially) and span (longest dependent path in parallel). The costs for a sample of the supported functions are given in Table 1. These costs can help the user decide on a block size for their particular application—a parameter that can be specified when creating a collection.

Using PaC-trees we have implemented CPAM: a C++ library which implements the full functionality of PAM, along with significant extra functionality involving compression. By default CPAM supports difference (or delta) encoding [43] within the blocked leaves. In such an encoding, each element is encoded based on the value of the previous element in the collection. This can greatly reduce space when elements that are close in the ordering of the collection are related. For example, if a graph is numbered so that neighboring vertices have similar indices, then the neighbors in a neighbor list will have small differences. These small numbers can then be encoded in a handful of bits each [49]. Similarly in an inverted index where each word points to a sequence of documents it appears in, if the documents are sorted, the differences between adjacent document identifiers can be small. This is especially true for common words, which take up the bulk of the space. In the paper we bound the extra space needed (due to the index using the tree structure) for PaC-trees compared to a static representation of the data (i.e., an array) directly using difference encoding (see Theorem 4.2).

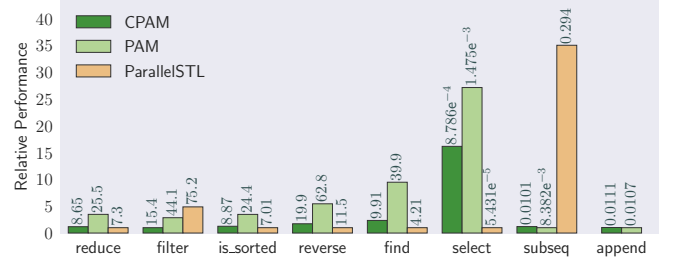


Figure 2. Relative performance of sequence primitives in CPAM (using $B = 128$), PAM, and ParallelSTL [33] on a 72-core machine with 2-way hyper-threading enabled. The numbers shown on top of the bars are the parallel (144-thread) running times in milliseconds. Lower is better. All benchmarks are run on sequences of length 10^8 containing 8-byte elements. For append, ParallelSTL takes 17.7 milliseconds on average (1594x larger than append in CPAM). CPAM and PAM represent sequences using purely-functional trees, whereas ParallelSTL uses arrays (hence static).

In our default blocked representation, the first element of a block is represented uncompressed, and the rest of the elements are compressed relative to the previous element. In addition to delta-encoding, CPAM also supplies an interface for the user to define their own form of compression for each block. For example, they can quantize values, or use other variable length codes when keys are known to be small. CPAM uses a reference counting garbage collector to manage the memory for both the internal nodes and the compressed leaf nodes, which can be of variable size due to compression.

CPAM supports augmentation in which each tree node maintains an aggregate of the values of its subtree (see more details in Section 3). The aggregation function is declared as part of the type of the tree. Augmentation is useful in many applications, and indeed we use it in all of the applications we describe later. PaC-trees store an augmented value per internal node, and one for each block at the leaves. Storing one value per block significantly reduces space relative to P-trees in PAM, which store a value for every element.

To demonstrate the effectiveness of PaC-trees, and their implementation in CPAM, we measure performance and space usage on (1) a collection of microbenchmarks that directly use some of the functions supported by the library, and (2) a handful of real-world applications.

For the microbenchmarks, we compare the performance of CPAM to PAM, and for sequences to the Intel implementation of the C++17 parallel STL library [33] (ParallelSTL). ParallelSTL is a highly optimized library supporting only sequences based on arrays. A summary of the results for sequences is given Fig. 2, and details including performance of ordered maps, and augmented maps are given in Section 10. Compared to PAM, CPAM achieves significantly better performance due to the reduced memory footprint, and hence reduced number of cache misses, while only requiring about 1/4-th as much space even without compression. Compared to ParallelSTL, CPAM has similar performance on operations

that visit the whole sequence, like reduce, but is significantly slower on n th since it requires $O(\log n + B)$ work as opposed to $O(1)$ for a random array access for ParallelSTL. On append CPAM is significantly faster since it requires $O(\log n + B)$ work to join to trees instead of $O(n)$ required by ParallelSTL to copy the input arrays into the output array.

We consider four applications: graphs, inverted indices, 2D range queries and 1D interval queries. For inverted indices, 2D range query and 1D interval query, CPAM achieves competitive performance to PAM while using 2.1x–7.8x less space. For graph processing, we compare to an existing system Aspen [25] that represents graphs using trees. CPAM uses 1.3–2.6x less space compared to Aspen, and is almost always faster than Aspen in all tested graph algorithms.

The main contributions of this paper are:

- A new functional data structure, PaC-trees, and associated parallel algorithms that support compression for sequences, sets, maps and augmented maps.
- Theoretical bounds on the costs (work and span) and the space of the data structure and associated algorithms.
- An implementation of PaC-trees as a library, CPAM, supporting the full functionality of PAM in addition to supporting default and user defined compression schemes.¹
- An experimental evaluation of the ideas and implementation on microbenchmarks and non-trivial applications.

2 Related Work

Our work extends P-trees and their C++ implementation in PAM [52]. Our key contribution is the ability to compress the trees achieving up to an order-of-magnitude reduction in space. This is achieved while being able to present cost bounds both in terms of time and space. These bounds are a function of a block size the user can select.

B-trees [6] and their variants block not just the leaves but all nodes of a tree, such that internal nodes can have a high fan-out. They are widely used in practice, especially for disk based data structures since nodes are on the scale of a page on disk and can be retrieved efficiently. However they are less relevant in the context of purely functional in-memory trees. In particular, path copying requires that an update copy all nodes on the path from the root to the leaf. If the nodes are large (e.g. 128+ elements each, as in our leaves) this copying would be very expensive both in terms of space and time. Various work has suggested blocking the leaves of a binary tree to represent sequences [1, 8, 16, 30, 38]. The idea is to reduce the cost of operations such as append or subsequence relative to array representations. As far as we know, these ideas have never been applied to ordered sets

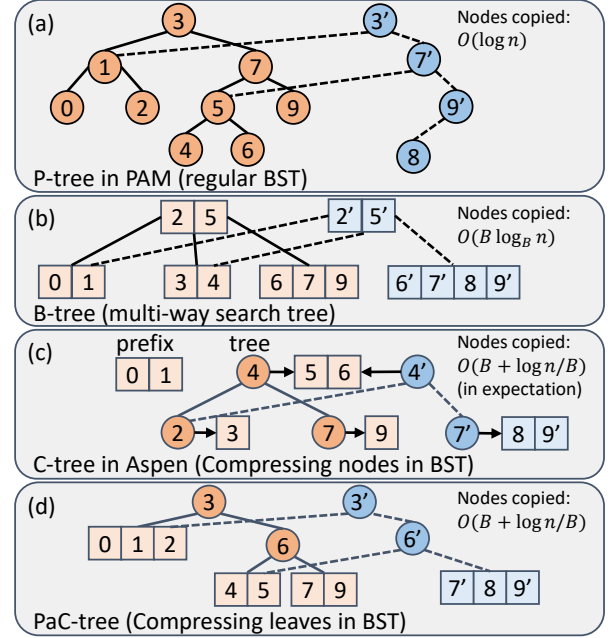


Figure 3. An illustration of (a) P-tree in PAM [11, 52] (regular BST), (b) B-tree (multi-way search tree), (c) C-tree [25] in Aspen (compressing all nodes in a BST) and (d) our PaC-tree (compressing all leaves in a BST) in CPAM. The orange nodes show a tree with keys 0-7 and 9. We then consider inserting a key 8. Blue nodes are what we need to create (copy or new) due to path-copying. Round nodes are tree nodes each storing a single key, and square nodes are organized in blocks of size $O(B)$ (expected for C-trees). Let n be the tree size, an insertion needs to copy $O(\log n)$ nodes in P-tree, $O(B \log_B n)$ in B-tree, and $O(B + \log(n/B))$ in C-tree (in expectation) or PaC-tree.

or ordered maps.² We also do not know of work that then compresses within the blocks.

Aspen [25] is a system for graph processing, based on purely functional trees and uses compression for the neighbor lists. At a high-level, our goals are shared with Aspen (e.g., non-mutating updates), but Aspen has several limitations. Importantly it is only designed for graphs, supporting only a small part of the functionality of CPAM. The tree representation in Aspen is also very different. It randomly selects elements from the collection to be *heads*. It then attaches a block of nodes to each head corresponding to the keys between the head and the next head, and puts the heads into a binary tree. PaC-trees do not require randomization, and have stronger theoretical bounds for primitive operations such as union than the bounds provided by C-trees in Aspen. We use CPAM to implement the full functionality of Aspen and compare to Aspen in Section 10.5.

¹We have made CPAM publicly available: <https://github.com/ParAlg/CPAM>.

²We note that the design of the chunked sequence datatype [1] could in principle be extended to support sets, maps, and augmented maps, although the implementation is specialized for ephemeral sequences.

Fig. 3 compares P-trees from PAM, functional B-trees, C-trees from Aspen, and PaC-trees. The comparison illustrates how they differ when inserting a new key.

Like CPAM, the Apache Spark [54] system supports a functional interface for collections. However it has significant differences. Firstly it only supports unordered sets. Secondly although it has a shared-memory parallel implementation, it is primarily designed for a distributed setting. This means its shared-memory implementation is not ideal.³

There is extensive research on concurrent tree data structures [3, 4, 18, 20, 28, 39, 44]. This work is mostly orthogonal to our work. Such trees support a fraction of the functionality of CPAM, typically just supporting linearizable inserts, deletes, updates and finds. Some recent work support range queries [5, 29], or arbitrary queries on a snapshot [53]. On the other hand concurrent trees support asynchronous updates, which PaC-trees do not—such updates are inherently non-functional. To support multiple concurrent updates, PaC-trees would require batching the update and applying as a batch in parallel (fairly comparing concurrent and batched structures like PaC-tree seems challenging for this reason). We expect the use cases would be quite different.

Blandford and Blelloch developed tree structures for ordered sets that support compression [9]. They present space bounds that are similar to ours, in terms of relating the space of a difference encoded sequence to the space of the data structure. However they support a small fraction of the functionality described in our work.

Functional trees using path-copying date back to at least the early 1990s [2], and in the sequential setting have been studied by Kaplan and Tarjan [36] and Okasaki [47].

3 Preliminaries

Binary search trees. A *binary search tree* (BST) is either an empty node, denoted as *nil*, or a node consisting of a *left* BST T_L , a key k (or with an associated value), and a *right* BST T_R , denoted $\text{node}(T_L, k, T_R)$, where k is larger than all keys in T_L and smaller than all keys in T_R . We use $lc(T)$ and $rc(T)$ to extract the left and right subtrees of T , respectively, and use $k(T)$ to denote the key stored at T 's root. The *size* of a BST T , or $|T|$, is the number of nodes in T . The *weight* of a BST T , or $w(T)$, is $1 + |T|$. The *height* of a BST T , or $h(T)$, is 0 for *nil*, and $\max(h(lc(T)), h(rc(T))) + 1$ otherwise. A tree node is a leaf if it has no children, and a regular node otherwise. The *left (right) spine* of a binary tree is the path of nodes from the root to a *nil* node, always following the left (right) tree.

A *weight-balanced tree*, or $\text{BB}[\alpha]$ trees [46] is a BST where for every $T = \text{node}(T_L, v, T_R)$, $\alpha \leq \frac{w(T_L)}{w(T)} \leq 1 - \alpha$. We omit

the parameter α with clear context. A weight-balanced tree T has height at most $\log_{\frac{1}{1-\alpha}} w(T)$.

Parallelism. Our implementation of PaC-trees is based on nested fork-join parallelism [23, 31, 35]. We analyze our algorithms use work-span model based on binary-forking [12]. The *work* W of a parallel algorithm is the total number of operations, while the *span* is the critical path length of its computational DAG. We use $s_1 \parallel s_2$ to indicate that statements s_1 and s_2 can run in parallel. Almost all algorithms use divide-and-conquer to enable parallelism. Any computation with W work and S span will run in time $T < \frac{W}{P} + S$ on P processors assuming shared memory and a greedy scheduler [15, 19]. We use $\log n$ to denote $\log_2(n+1)$ in the cost bounds.

Encoding schemes. We use *Difference Encoding* (DE) to encode integer keys. Given a sorted set of keys, K , the difference encoding scheme stores the differences between consecutive keys using an integer code, such as byte or γ codes. We only consider byte codes in this paper since they are cheap to encode and decode and do not waste much space compared to using γ codes [49].

Functional data structures. PaC-trees are purely functional data structures. In functional data structures values are immutable, so updates must be made by copying parts of the structure. For search trees, only the path to the update location needs to be copied. Hence for balanced trees of size n , single point updates such as inserts and deletes involve copying $O(\log n)$ nodes (Fig. 3(a)). This also applies to multi-point updates. For example, if a filter ends up removing a single element, only $O(\log n)$ nodes need to be copied. Functional trees can also easily support multiversioning with low time and space overhead [7, 51]. Because the data are immutable, any operation accesses the tree in an isolated version. Updates can be applied in *batches* in parallel and yield a new version. This enables all read-only queries to be performed at the same time without being affected by ongoing (concurrent) updates. In addition to multiversioning, functional data structures also allow for multiple histories.

Join-based algorithms. PaC-trees are implemented using the *join-based* approach [11, 14, 32, 50–52] first implemented in PAM [52]. In the framework, a variety of tree algorithms are implemented based on two primitives, join and expose.⁴ Given a balancing scheme \mathcal{S} , the $\text{join}(T_L, e, T_R)$ function returns a balanced tree T satisfying \mathcal{S} which has the same in-order values as $\text{node}(T_L, e, T_R)$. In other words, it concatenates T_L and T_R by an entry e in the middle while preserving the balancing invariants (see Fig. 7 as an example of joining two PaC-trees). The $\text{expose}(T_L)$ function returns a triple (T_L, e, T_R) , where $e \in T$ is an entry, T_L and T_R are two binary trees such that both T_L and T_R satisfy \mathcal{S} , are balanced with each other under \mathcal{S} , and T_L (T_R) contains all keys in

³Their shared-memory implementation is between 3.2–4.9x slower than CPAM for a map, reduce, and group-by style example taken from their user guide. For primitives such as map and reduce, their implementation performs up to 2 orders of magnitude worse than CPAM (see Section 10.2)

⁴PAM did not explicitly use expose as a primitive, but only conceptually treated it as a primitive.

T that go before (after) e in T 's in-order value. It has been shown that on weight-balance trees with $\alpha \leq 1 - 1/\sqrt{2}$, a join operation can be done in $O(\log \frac{n}{m})$ work [11], where $n = \max(|T_L|, |T_R|)$ and $m = \min(|T_L|, |T_R|)$.

Based on join and expose, many parallel tree algorithms can be expressed in a simple and elegant recursive style (see Fig. 5, Fig. 6, and Fig. 10 for examples). We adopt the join-based approach in our implementation, and in particular carefully designed join and expose functions for PaC-trees. This greatly simplifies the implementation and correctness arguments of our algorithms. We give more details in Sections 5 and 6.

Augmentation. An *augmented tree* is a search tree where each node maintains an aggregated value (called **augmented values**) of all entries in its subtree. Typical examples would be a weighted sum, minimum or maximum of values, where we can obtain the augmented value in a node by combining augmented values of the children and itself. This generalizes to all associative operations. PaC-trees support generic user-defined augmentation for any associative operations. An example of PaC-tree with augmentation is shown in Fig. 4.

4 PaC-Trees

In this paper, we propose PaC-trees to support purely functional collections, which support *parallelism*, *determinism*, *compression*, *augmentation*, *strong theoretical bounds*, and *multi-versioning*. PaC-trees are purely functional. The base data structure of a PaC-tree is a weight-balanced BST. The internal nodes remain binary so they are cheap to copy. The leaves in a PaC-tree are organized in blocks of size B to $2B$ for some parameter B . An illustration is shown in Fig. 3. If the blocks grow too large, they are split, and if they become too small they are merged with a neighboring node.

Definition 4.1 (PaC-tree). A PaC-tree $\text{PaC}(\alpha, B, C)$, parameterized by the balancing factor α , block size B , and encoding scheme C satisfies the following invariants:

- **(Weight Balance)** For any tree node $v \in T$, $\alpha \leq \frac{w(v_*)}{w(v)} \leq 1 - \alpha$, where $\alpha \leq 1 - \frac{1}{\sqrt{2}}$ is a constant, and v_* is either $lc(v)$ or $rc(v)$. Unless mentioned otherwise, we use $\alpha = 0.29$.
- **(Blocked Leaves)** If $|T| \geq B$, each leaf $u \in T$ maintains B to $2B$ entries in an array (called a **block**) using the encoding scheme C . Unless mentioned otherwise, we assume C is empty, which means the entries are blocked without additional compression of the entries.

When the context is clear, we omit α , B and C in the definition and simply call it a **PaC-tree**. We call a leaf node containing multiple entries in a PaC-tree a **flat node**, and a node containing a single entry a **regular node**. We say a PaC-tree (or a subtree) T is a **simplex** tree if $|T| < B$, and thus T only contains regular nodes. We say a PaC-tree (or a subtree) T is a **complex** tree if T contains both regular nodes and flat nodes. We define the **expanded** version of

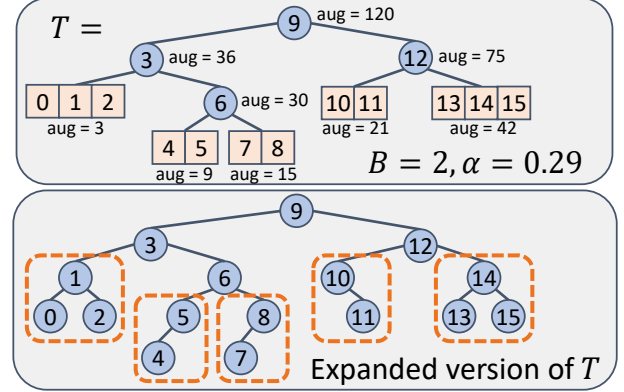


Figure 4. (a). An illustration of a PaC-tree with keys $\{0, 1, \dots, 15\}$, and augmentation as sum of keys. All nodes are weight-balanced. All leaves are blocked as arrays of size B to $2B$. (b) The expanded version of the PaC-tree in (a).

a PaC-tree T (or a flat node v) to be a regular binary tree (without flat nodes), where all flat nodes in T (or v itself) are fully expanded as perfectly-balanced binary trees. In Fig. 4, we show an example of an expanded tree.

We now present the space bound of a PaC-tree. For integer keys, we can use difference encoding to bound the space.

Theorem 4.2. *The total space of a PaC-tree $\text{PaC}(\alpha, B, C_{DE})$ maintaining a set E of integer keys is $s(E) + O(|E|/B + B)$, where C_{DE} is difference encoding, and $s(E)$ is the size needed for E using difference encoding.*

Proof. The space needed for a PaC-tree includes the regular nodes and the leaf nodes. First of all, when $|E| < B$, all entries are maintained in a simplex tree, taking $O(B)$ space. When $|E| \geq B$, there are $O(|E|/B)$ regular nodes, each taking $O(1)$ space for meta-data (pointers, size, etc.). The total space used by regular nodes is $O(|E|/B)$. All the leaf nodes are organized in blocks. Let A be an array that stores all keys in E using difference encoding. Comparing the total size of all the blocks and A , the only extra space is the first element of each block (which cannot be compressed). There are $O(|E|/B)$ such blocks, and thus the extra space used is $O(|E|/B)$. \square

We note that this bound is deterministic, as opposed to the bound for C-trees (which only holds in expectation). Furthermore, using known facts about difference encoding yields the following result, showing that PaC-trees yield a compact parallel representation of ordered sets [9].

Corollary 4.3. *Given any set from $U = \{0, \dots, m-1\}$ with $|S| = n$, the total space of a PaC-tree $\text{PaC}(\alpha, B, C_{DE})$ maintaining S is $O(n \log \frac{n+m}{n})$ bits for $B = \Omega(\log n)$.*

5 Algorithms

We now describe join-based algorithms on PaC-trees. To enable a general ordered map interface, we implement PaC-trees based on the PAM interface. PAM supports dozens of

```

1  fold(T) {
2    flatten T into array A
3    (encoding if needed)
4    return A; }
5  unfold(A) {
6    /* return a perfectly balanced tree
7     from sorted array A */
8  expose(T) {
9    if (isflat(T)) {
10     T' = unfold(T);
11     return (lc(T'), k(T'), rc(T')); }
12   else return (lc(T), k(T), rc(T)); }
13  join(TL, k, TR) {
14    if (heavy(TL, TR))
15     return join_right(TL, k, TR);
16    if (heavy(TR, TL))
17     return join_left(TL, k, TR);
18    return node(TL, k, TR); }
19  /* join_left is symmetric */
20  join_right(TL, k, TR) {
21    (l, k', c) = expose(TL);
22    if (balance(|TL|, |TR|))
23     return node(TL, k, TR);
24    T' = join_right(c, k, TR);
25    (l1, k1, r1) = expose(T');
26    if (balance(|l|, |T'|))
27     return node(l, k', T');
28    if ((balanced(|l|, |l1|) and
29         balanced(|l| + |l1|, r1)))
30     return rotateleft(node(l, k', T'));
31    else return rotateleft(node(l, k',
32                               rotateright(T'))); }
33  join2(TL, TR) {
34    if (TL = nil) return TR;
35    (T'L, m, _) = split(TL, last(TL));
36    return join(T'L, m, TR); }
37  node(l, k, r) {
38    /* create node x with left subtree l,
39     root key k and right subtree r */
40    if (|x| > 4B) return x;
41    if (B ≤ |x| ≤ 2B) return fold(x);
42    else { // 2B < |x| ≤ 4B
43      /* redistribute x's both subtrees to
44       be flat nodes with |x|/2 entries */
45      return x; } }
46  split(T, k) {
47    if (|T| = 0) return (nil, nil, nil);
48    (L, m, R) = expose(T);
49    if (k == k(m)) return (L, m, R);
50    if (k < k(m)) {
51      (LL, b, LR) = split(L, k);
52      return (LL, b, join(LR, m, R));
53    } else {
54      (RL, b, RR) = split(R, k);
55      return (join(L, m, RL), b, RR); } }

```

Figure 5. Primitives on PaC-trees. All codes are functional (e.g. rotates copy nodes).

```

1  from_sorted(A, n) {
2    if (n = 0) return nil;
3    if (n = 1) return node(nil, A[0], nil);
4    L = from_sorted(A, n/2) ||
5    R = from_sorted(A+n/2, n-n/2);
6    return node(L, A[n/2], R); }
7  build(A, n) {
8    parallel_sort(A, n);
9    return from_sorted(A, n); }
10 // keep a key in T only when it satisfies f
11 filter(T, f) {
12   if (T == nil) return nil;
13   (L, k, R) = expose(T);
14   TL = filter(L, f) ||
15   TR = filter(R, f);
16   if (f(k))
17     return join(TL, k, TR);
18   else return join2(TL, TR); }

```

Figure 6. Examples of parallel algorithms on PaC-trees. “||” indicates calls that are made in parallel.

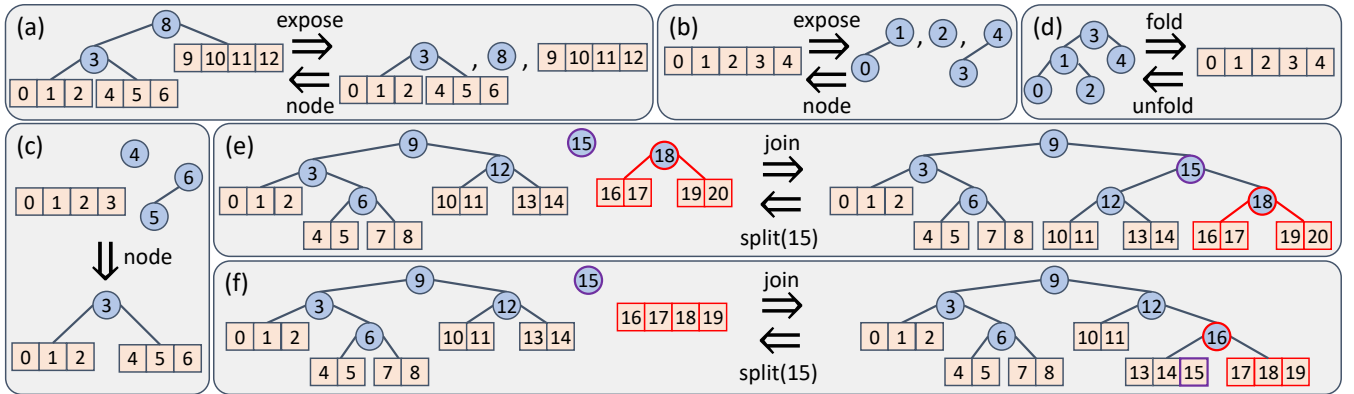


Figure 7. Illustration of primitives on PaC-trees. For Figures (a)–(d), $B = 3$. For Figures (e)–(f), $B = 2$. Fig. (a): the expose function on a regular node and the node function to obtain a regular node when the output tree size is larger than $4B$. Fig. (b): the expose function on a flat node and the node function to obtain a flat node when the output tree weight is between B and $2B$. Fig. (c): the node function to obtain a flat node when the output size is between $2B$ and $4B$. Fig. (d): fold and unfold functions. Fig. (e): join function on two regular nodes and its corresponding split function. Fig. (f): join function on a regular node and a flat node and its corresponding split function.

operations on sequences, sets, maps, and augmented maps, and it would require significant work to re-implement them all. Instead, we carefully redesigned join and expose such

that *all the other algorithms can remain the same as in PAM*. In particular, none of the other algorithms have to deal with the blocked leaves or compression, which greatly simplifies

the algorithm design and correctness arguments. We found that the overhead of this approach is not large, but for many frequently-used operations, we design special base cases for dealing with compressed nodes. These base cases can improve the performance by up to 6x (see Section 8). Some of the theoretical results also require special base cases (see Section 6.1).

At a high-level, when exposing a flat node, the node is automatically expanded (using `unfold`), and similarly when join obtains a complex tree of size B to $2B$, it is flattened (fold). An illustration of `unfold` and `fold` is shown in Fig. 7. We start with the `join` and `expose` algorithms. We then present the `union` algorithm as an example to illustrate join-based algorithms, and give the code for other functions in Fig. 6 and Fig. 10. We focus on `union` as it is the core subroutine used in applications such as inserting or deleting batches of vertices and edges in graphs, combining inner trees when constructing range trees, and updating sets of documents in an inverted index, among others.

Expose. This function returns the left subtree, root data and the right subtree of a node T . For a regular node, this function just reads the child pointers and the root. For a flat node, this function first unfolds the tree into a perfectly balanced tree and then reads the corresponding data.

Join. Recall that the `join` function takes two trees T_L and T_R , and a key k (or a key-value) as input, and returns a balanced tree concatenating entries in T_L , k and T_R in order (see Fig. 7). In other words, when trees are used for ordered sets or maps, k should be larger than all keys in T_L and smaller than all keys in T_R . Pseudocode for `join` is shown in Fig. 5.

The algorithm first compares the weights of T_L and T_R . When balanced, they are directly connected by k . The other two cases are symmetric so WLOG we assume $|T_L| > |T_R|$. In this case, the algorithm must attach T_R in the right spine of T_L , which will be handled by `join_right`(T_L, k, T_R). This algorithm first checks if T_L and T_R are balanced and connects them if so. Otherwise, it recursively calls `join_right` on $rc(T_L)$ and T_R , getting T' . If we re-attach T' as T_L 's right child, we will get a “correct” output tree (modulo balance). We then use a single or double rotation to rebalance if necessary. It is known that either a single or double rotation can rebalance a weight-balanced tree in this situation [11]. This guarantees the *weight balance* invariant of PaC-trees.

To also guarantee the *blocked leaves* invariant, we add two conditions when calling `node` to create a new node with its left and right subtrees. Whenever a node with size B to $2B$ is created, we fold the tree into a flat node. Whenever a node with size $2B$ to $4B$ is created, we extract the median of the tree as the root to re-distribute its two subtrees, such that both subtrees are flat nodes with (almost) the same size.

Lemma 5.1. *The join function maintains the invariants of PaC-trees.*

Split. For a PaC-tree T and key k , `split`(T, k) returns a triple (T_L, b, T_R) , where T_L (T_R) is a tree containing all keys in T that are less (greater) than k , and b the entry of key k if $k \in T$ (see Fig. 7). We first use `expose`(T) to get its left (right) subtrees $lc(T)$ ($rc(T)$) and root key $k(T)$, and compare k with $k(T)$. If $k = k(T)$, we simply return $(lc(T), k, rc(T))$. Otherwise WLOG we assume k is smaller. In that case, the entire right subtree $rc(T)$ and the root $k(T)$ belong to T_R . We then split $lc(T)$ by k , getting (L_L, b, L_R) . By definition, all keys smaller than k should be in L_L , and all keys larger than k can be obtained by `join`($L_R, k(T), rc(T)$).

Union. Using `join` and `split`, we can implement set algorithms on two PaC-trees, such as `union`, `intersection` and `difference`. We describe `union` as an example (the other two are similar). This algorithm uses divide-and-conquer. At each level of recursion, T_1 is split by the root of T_2 , breaking T_1 into two subsets with all keys smaller (larger) than $k(T_2)$, denoted as L_1 (R_1). Then two recursive calls to `union` are made in parallel. One unions $L(T_2)$ with L_1 (all keys smaller than $k(T_2)$), returning T_L , and the other one unions $R(T_2)$ with R_1 (all keys larger than $k(T_2)$), returning T_R . Finally the algorithm combines the results with `join`($T_L, k(T_2), T_R$).

Other algorithms. We present the pseudocode for the other two set algorithms (`intersection` and `difference`) in Fig. 10. We also show the code for three other useful functions, `multi_insert`, `map` and `reduce` in Fig. 8. We note that these algorithms are exactly the same as in PAM, by extracting out the semantics of `join` and `expose`.

Importantly, all of our PaC-tree algorithms are theoretically efficient. We present the work-span bound in Table 1 and give a proof for `union` as an example in Section 6. Note that Lemma 5.1 ensures the correctness of the other algorithms, as their return values are always obtained by a `join`.

Theorem 5.2. *All join-based algorithms on PaC-tree maintains the invariants of PaC-trees.*

6 Theoretical Guarantees

In the following section we show work and span bounds for operations on PaC-trees. We assume the encoding scheme is empty, which means that to flatten or expand a block of size n costs $O(n)$ work and $O(\log n)$ span. If the encoding scheme is not parallelizable (e.g., for difference encoding), the span bound of the algorithms will be affected. We present more details in Section 6.2.

We start with the cost of the `join` and `split` algorithms.

Theorem 6.1. *Consider a join algorithm on two PaC-trees T_L, T_R and an key k . Let $n = \max(|T_L|, |T_R|)$ and $m = \min(|T_L|, |T_R|)$. If both T_L and T_R are complex trees, the algorithm takes $O(\log \frac{n}{m})$ work and span. If both T_L and T_R are simplex trees, the algorithm takes $O(B)$ work and $O(\log B)$ span. Otherwise, the algorithm takes $O(B + n/B)$ work and $O(\log n)$ span.*

```

1  m_ins_helper( $T, A, m$ ) {
2    if ( $T == \text{nil}$ ) return from_sorted( $A, m$ );
3    if ( $n == 0$ ) return  $T$ ;
4    ( $L, k, R$ ) = expose( $T$ );
5     $s = \text{binary\_search}(A, m, k)$ ;
6    if ( $A[s] == k$ )  $b = 1$ ;
7     $T_L = \text{m\_ins\_helper}(L, A, s)$ ;  $||$ 
8     $T_R = \text{m\_ins\_helper}(R, A+s+b, m-s-b)$ ;
9    return join( $T_L, k, T_R$ ); }

10 multi_insert( $T, A$ ) {
11    $A' = \text{parallel\_sort}(A)$ ;
12   return m_ins_helper( $T, A', |A|$ ); }

13 //  $f$  is an associative binary operation
14 //  $I$  is the identity
15 reduce( $T, f, I$ ) {
16   if ( $T$  is nil) return  $I$ ;
17   ( $L, k, R$ ) = expose( $T$ );
18    $x = \text{reduce}(L, f, I)$ 
19    $|| y = \text{reduce}(R, f, I)$ ;
20   return  $f(f(x, k), y)$ ; }

21 // map each entry in  $T$  using function  $f$ 
22 map( $T, f$ ) {
23   if ( $T == \text{nil}$ ) return nil;
24   ( $L, k, R$ ) = expose( $T$ );
25    $T_L = \text{map}(L, f)$ 
26    $|| T_R = \text{map}(R, f)$ ;
27   return join( $T_L, k, T_R$ ); }

```

Figure 8. Some other algorithms on PaC-trees.

```

1  join( $T_L, k, T_R$ , expand=false) {
2    if (heavy( $T_L, T_R$ )) return join_right( $T_L, k, T_R$ , expand);
3    if (heavy( $T_R, T_L$ )) return join_left( $T_L, k, T_R$ , expand);
4    return node( $T_L, k, T_R$ ); }

5  join_right( $T_L, k, T_R$ , expand=false) {
6    ( $l, k', c$ ) = expose( $T_L$ );
7    if (balance( $|T_L|, |T_R|$ )) return node( $T_L, k, T_R$ , expand);
8     $T' = \text{join\_right}(c, k, T_R)$ ;
9    ( $l_1, k_1, r_1$ ) = expose( $T'$ );
10   if (balanced( $|l|, |T'|$ )) return node( $l, k', T'$ , expand);
11   if ((balanced( $|l|, |l_1|$ )) and (balanced( $|l| + |l_1|, r_1$ )))
12     rotate_left(node( $l, k', T'$ ));
13   else rotate_left(node( $l, k', \text{rotate\_right}(T')$ )); }

14 join_left( $T_L, k, T_R$ ) { /* symmetric to join_right */ }

15 node( $l, k, r$ , expand=false) {
16   let the left child of  $k$  be  $l$ ;
17   let the right child of  $k$  be  $r$ ;
18   if (expand) return  $k$ ;
19   if ( $|k| > 4B$ ) return  $k$ ;
20   if ( $B \leq |k| \leq 2B$ ) return fold( $k$ );
21   else { //  $2B < |k| \leq 4B$ 
22     /* redistribute  $k$ 's both subtrees to be flat nodes with  $|k|/2$  entries */
23     return  $k$ ; } }

24 expose( $T$ ) {
25   if (isregular( $T$ )) return ( $lc(T), k(T), rc(T)$ );
26   else {
27      $T' = \text{unfold}(T)$ ;
28     return ( $lc(T'), k(T'), rc(T')$ ); } }

29 fold( $T$ ) {
30   if ( $B \leq w(T) \leq 2B$ ) {
31     flatten  $T$  into array  $A$ 
32     (encoding if needed)
33     return  $A$ ;
34   } else return  $T$ ; }

35 unfold( $A$ ) {
36   /* build a perfectly balanced tree  $T$ 
37   from entries in (sorted) array  $A$  */ }

38 refold( $T$ ) {
39   if ( $T$  is not marked) return  $T$ ;
40   if ( $B \leq |T| \leq 2B$ ) return fold( $T$ );
41   else {
42     ( $L, m, R$ ) = expose( $T$ );
43      $T_L = \text{refold}(L)$   $|| T_R = \text{refold}(R)$ ;
44     return join( $T_L, m, T_R$ ); } }

```

Figure 9. Some useful primitives for join-based algorithms on PaC-trees for more efficient set algorithms.

Proof. WLOG, let's assume $n = |T_L| \geq |T_R| = m$.

For two complex trees, we first prove that join never decompresses a leaf. Note that the algorithm will follow the right spine of the tree until finding a subtree t in T_L that is balanced with T_R , we will prove that there exist a regular node t that is balanced with T_R . This is because as a complex tree, T_R has size at least $2B$. Along the right spine of T_L , the smallest complex subtree has size at most $4B$. For any $\alpha \leq 1/3$, we must be able to find a complex subtree in T_L that is balanced with T_R . This proves that the total number of tree nodes we need to visit on the right spine is $O(\log \frac{n}{m})$.

For two simplex trees, the work is no more than copying both T_L and T_R and concatenating them, which is $O(B)$ work and $O(\log B)$ span.

If T_L is a complex tree and T_R is a simplex tree, we need to first follow the right spine to find a leaf node l in T_L , which takes $O(\log \frac{n}{B})$ time. Then it combines the leaf node with T_R , which flattens both l and T_R , concatenates them, and rebalance the result. This process takes no more than $O(B)$ work and $O(\log B)$ span. \square

Theorem 6.2. Consider a split algorithm on a PaC-tree T . If T is a complex tree, the work and span of split are

```

1  splitLast(T) {
2    (L, k, R) = expose(T);
3    if (R == nil) return (L, k);
4    else {
5      (T', k') = splitLast(R);
6      return (join(L, k, T'), k'); } }

```

```

7  join2(T_L, T_R) {
8    if (T_L == nil) return T_R;
9    else {
10     (T'_L, k) = splitLast(T_L);
11     return join(T'_L, k, T_R); } }

```

Simple version used in implementation:

```

12 union(T1, T2) {
13   if (T1 == nil) return T2;
14   if (T2 == nil) return T1;
15   (L2, k2, R2) = expose(T2);
16   (L1, b, R1) = split(T1, k2);
17   T_L = union(L1, L2);
18   || T_R = union(R1, R2);
19   return join(T_L, k2, T_R); }

```

```

20 intersect(T1, T2) {
21   if (T1 == nil) return nil;
22   if (T2 == nil) return nil;
23   (L2, k2, R2) = expose(T2);
24   (L1, b, R1) = split(T1, k2);
25   T_L = intersect(L1, L2);
26   || T_R = intersect(R1, R2);
27   if (b) return join(T_L, k2, T_R);
28   else return join2(T_L, T_R); }

```

```

29 difference(T1, T2) {
30   if (T1 == nil) return nil;
31   if (T2 == nil) return T1;
32   (L2, k2, R2) = expose(T2);
33   (L1, b, R1) = split(T1, k2);
34   T_L = difference(L1, L2);
35   || T_R = difference(R1, R2);
36   return join2(T_L, T_R); }

```

Special base cases for tighter bound:

```

37 union_(T1, T2) {
38   if (T1 == nil) return T2;
39   if (T2 == nil) return T1;
40   if (isflat(T1) or isflat(T2)) {
41     return refold(union_base(T1, T2)); }
42   (L2, k2, R2) = expose(T2);
43   (L1, b, R1) = split(T1, k2);
44   T_L = union_(L1, L2) || T_R = union_(R1, R2);
45   return join(T_L, k2, T_R); }

```

```

46 union_base(T1, T2) {
47   if (T1 == nil) return T2;
48   if (T2 == nil) return T1;
49   if (isflat(T1)) unfold(T1);
50   if (isflat(T2)) unfold(T2);
51   (L2, k2, R2) = expose(T2);
52   (L1, b, R1) = split(T1, k2, true);
53   T_L = union_base(L1, L2) || T_R = union_base(R1, R2);
54   return join(T_L, k2, T_R, true); }

```

Figure 10. Set algorithms on PaC-trees.

$O(\log \frac{|T|}{B} + B)$ and $O(\log |T|)$, respectively. If T is an simplex tree, the work and span of split is $O(\log |T|)$.

Proof. For a simplex tree the cost directly follows the result on P-trees [11]. For a complex tree, the only difference of split on PaC-trees from split on PAM trees is the unfold performed in expose and fold in join. It takes $O(\log |T|/B)$ steps to reach a leaf node. The unfold function in expose (Line 48) is performed at most once. For all join calls, at most one on each side can involve a simplex tree. So the total overhead is at most $O(B)$ in work and $O(\log B)$ in span. \square

Based on these results, we now analyze the cost of the set operations.

Theorem 6.3. *Consider the union algorithm (and the closely related intersection and difference algorithms) in Fig. 10 on two PaC-trees of sizes m and $n \geq m$. The work and span for these algorithms are $O(m \log \frac{n}{m} + mB)$ and $O(\log n \log m)$ respectively.*

To prove the theorem, we first present some definitions and lemmas. First, note that all the work can be asymptotically bounded by the three categories below:

- (1). **split work**: all work done by split (Line 16),
- (2). **join work**: all work done by join (Line 19) or join2 in intersection and difference,
- (3). **expose work**: all work done by expose (Line 15).

One observation is that the split work is identical among the three set algorithms. This is because the three algorithms behave the same on the way down the recursion when doing splits, and only differ in what they do at the base case and on the way up the recursion when building the output tree (see the other two set algorithms in Fig. 10).

We use op to denote the set operation (one of union, intersection or difference). In these algorithms, the tree T_1 is split by the keys in T_2 . We call T_1 the **decomposed tree** and T_2 the **pivot tree**, denoted as T_d and T_p respectively. Let $m = \min(|T_p|, |T_d|)$ and $n = \max(|T_p|, |T_d|)$.

Lemma 6.4. *For each function call to op on trees $P \subseteq T_p$ and $D \subseteq T_d$, the work done by join (or join2) is asymptotically bounded by the work done by split.*

Proof. Assume the return value is R .

First of all, the work of split is $\Theta(\log |D| + B)$. Note that the work of join (or join2) can be bounded by $O(\log |R| + B)$.

Notice that difference returns the keys in $D \setminus P$. Thus for both intersection and difference we have $R \subseteq D$. Therefore $|R| \leq |D|$, which means the work done by join or join2 is no more than the work done by split.

For union, first of all, we always call join instead of join2. If $|P| \leq |D|$, then $|R| \leq 2|D|$. join costs work $O(\log |R| + B) = O(\log |P| + B)$, which is no more than $\Theta(\log |D| + B)$.

Consider $|P| > |D|$. The subtrees $lc(P)$ and $rc(P)$, which are used in the recursive calls, have size at least $\alpha|P|$ and at most $(1 - \alpha)|P|$. After combining with a subset of elements in D (which has size smaller than $|P|$), the return value of each recursive call should have size at least $\alpha|P|$ and $(2 - \alpha)|P|$. Denote these two trees from recursive calls as t_l and t_r , respectively. Note that α is a constant, so the difference of size between t_l and t_r is also no more than a constant. WLOG assume $|t_l| \geq |t_r|$. In the following, we discuss different cases of whether t_l and t_r are complex or simplex trees. We will show that, in all cases, joining $|t_l|$ and $|t_r|$ has work $O(\log |D| + B)$.

1. When both t_l and t_r are simplex trees. From Theorem 6.1, join costs $O(B)$ work.
2. When both t_l and t_r are complex trees. From Theorem 6.1, join costs $O(\log \frac{|t_l|}{|t_r|}) = O(1)$ work.
3. When t_l is a complex tree, but t_r is a simplex tree. This means that $|t_l| > B$ and $|t_r| \leq B$. From Theorem 6.1, join costs $O\left(B + \log \frac{|t_l|}{B}\right)$ work. Note that, as stated above, $|t_r| \geq \alpha|P|$. Considering $|t_r| \leq B$, we know that $|P| = O(B)$, which also indicates $|t_l| = O(B)$. Plug this into the work of join $O\left(B + \log \frac{|t_l|}{B}\right)$, we can get the bound $O(B)$.

In summary, in all cases the work of join or join2 is asymptotically bounded by the corresponding split function. \square

Next, we prove the bounds for split work and expose work, respectively.

Lemma 6.5. *The expose work is $O(\min(mB, n))$.*

Proof. expose costs $\Theta(B)$ when the subtree is a flat node, and $O(1)$ otherwise. At most $O(m)$ nodes in T_p will split T_d , so the total cost is $O(mB)$. The cost is also no more than $O(n)$ since each node is involved in at most one expose, after which the flat node will be fully expanded. In summary the cost is $O(\min(mB, n))$. \square

Lemma 6.6. *The total split work is $O(m \log \frac{n}{m} + mB)$.*

Proof. The total split work can be viewed as two parts: the total work to done by split functions to traverse and split non-flat nodes, and the work to expose and split the flat nodes. Note that here “non-flat nodes” include both regular nodes in complex trees, and all the nodes in expanded trees.

First of all, the total work to traverse and split all non-flat nodes can be asymptotically bounded by the split work when

both T_p and T_d are considered to be fully expanded. This cost is $O(m \log \frac{n}{m})$ from the result for P-trees [11].

We then consider all work done by split functions on flat nodes. The only extra cost is the cost of unfold. Every node in T_p will be used at most once to split T_d , which involves at most one unfold function with cost $O(B)$. There can be at most $O(m)$ nodes in T_p used to split T_d . Thus the total unfold work in split is $O(mB)$.

Therefore in total the split work is $O(m \log \frac{n}{m} + mB)$. \square

We can now prove Theorem 6.3.

Proof. (Theorem 6.3) Combining Lemmas 6.4 to 6.6 proves the work bound in Theorem 6.3. For the span, note that the algorithms need $O(\log |T_p|/B)$ rounds to reach a flat node, where the flat node will be expanded, taking $O(\log B)$ span. Then the algorithm keeps recursing until a *nil* node is reached, which takes $O(\log B)$ rounds. In each of the recursive calls, we need $O(\log |T_d|)$ span to deal with split and join. In total the span is $O(\log m \log n)$. \square

6.1 Set Algorithms with Better Work Bound

Note that the $O(mB)$ term can be expensive when m is large. In fact, we can show a tighter bound using a more efficient (but more complicated) base case, which we present next. We note that in our implementation, we use the version in Fig. 5, which has good performance in practice. The main result in this section is the theorem below, based on the algorithm shown in Fig. 10 as `union_`.

Theorem 6.7. *There exist algorithms for union, intersection and difference on two PaC-trees of sizes m and n ($n \geq m$) with work $O\left(m \log \frac{n}{m} + \min(n, mB)\right)$ and span $O(\log n \log m)$.*

Algorithm. The general idea is to avoid folding and unfolding simplex trees during the union algorithm. In particular, we hope each flat node is folded and unfolded $O(1)$ times during the entire union algorithm. To ensure this, we implemented a special base case for the set algorithms. We show the code for union in Fig. 10, the other two are similar. The base case will explicitly determine if the current input is a flat node. If any of them is, it will be expanded directly. In the subsequent join, we will pass an extra parameter to indicate that the tree is already expanded, and thus there is no need to fold or unfold them again in this join algorithm. This parameter is the last parameter of the join algorithm in Fig. 9. It is set to *false* by default, which makes it exactly the same as the version in Fig. 5. When it is set to *true*, the join algorithm will never fold any node. Instead, at the end of the base case, the entire result tree will be fixed using `refold` (Line 41), which traverses the tree and folds any subtree of size B to $2B$ back to blocks.

Theoretical Cost. As the span bound is not affected, we will only show the new proof for the work here. First, note that

total work for each of these algorithms can be considered as several parts:

- (1). all work done by unfold operations in the base cases, including possibly those in `split` and `join` (or `join2`) function calls, denoted as **unfold work**,
- (2). all work done by refold operations, denoted as **refold work**,
- (3). all work done by split operations except for the unfold work in base cases (already charged in (1).), denoted as **split work**,
- (4). all work done by join or join2 operations, denoted as **join work**.

We note that all the rest of the cost can be asymptotically bounded by the above four categories of work.

One observation is that the split work and unfold work are identical among the three set algorithms. This is because the three algorithms behave the same on the way down the recursion when doing splits, and algorithms only differ in what they do at the base case and on the way up the recursion when they join back.

We start with some notation. We follow some the notation used in [11]. Throughout the section, we use `op` to denote the algorithm or function call on union, intersection or difference, and use `op_base` to denote the corresponding base case algorithm or function call (using the unfolded version). In these three algorithms, the first tree (T_1) is split by the keys in the second tree (T_2). We call T_1 the **decomposed tree** and T_2 the **pivot tree**, denoted as T_d and T_p respectively. The returned tree of the algorithms is denoted as T_r . We use $m = \min(|T_p|, |T_d|)$ and $n = \max(|T_p|, |T_d|)$. We denote the subtree rooted at $v \in T_p$ as $T_p(v)$, and the tree of keys from T_d that v is operated with as $T_d(v)$ (i.e., `op`($v, T_d(v)$) or `op_base`($v, T_d(v)$) is called at some point in the algorithm. This essentially means that v 's subtree in T_p is processed with the tree $T_d(v)$ in a recursive call. Note that $T_d(v)$ may not be a subtree in T_d , but is a tree of keys as a subset of T_d . We call such $T_d(v)$ a **subset tree** of the decomposed tree T_d . We say $v \in T_p$ **processes** $T_d(v)$. For $v \in T_p$, we refer to $|T_d(v)|$ as its **splitting size**.

Lemma 6.8. *The refold work can be asymptotically bounded by the unfold work.*

Proof. We note that during the process of tracking down the tree, we will refold the subtree if and only at least a subset of it was previously unfolded at some point in this algorithm. Since refold costs $O(B)$ work, it can be asymptotically bounded by the corresponding unfold function invoked previously. \square

Theorem 6.9. *For each function call to `op` on trees $T_p(v)$ and $T_d(v)$, the work done the join (or join2) is asymptotically bounded by the work done by split.*

Proof. In the following, we use P and D to denote $T_p(v)$ and $T_d(v)$, respectively, for simplicity. Assume the return value is R . First of all, the work of split is $\Theta(\log |D| + B)$.

For intersection or difference, the work of join (or join2) is $O(\log |R| + B)$. Notice that difference returns the keys in $D \setminus P$. Thus for both intersection and difference we have $R \subseteq D$. Therefore $|R| \leq |D|$, which means the work done by join or join2 is no more than the work done by split.

For union, first of all, we always call join instead of join2. If $|P| \leq |D|$, then $|R| \leq 2|D|$. join costs work $O(\log |R| + B) = O(\log |P| + B)$, which is no more than $\Theta(\log |D| + B)$.

Consider $|P| > |D|$. The subtrees $lc(P)$ and $rc(P)$, which are used in the recursive calls, have size at least $\alpha|P|$ and at most $(1 - \alpha)|P|$. After combining with a subset of elements in D (which has size smaller than $|P|$), the return value of each recursive call should have size at least $\alpha|P|$ and $(2 - \alpha)|P|$. Denote these two trees from recursive calls as t_l and t_r , respectively. Note that α is a constant, so the difference of size between t_l and t_r is also a constant. WLOG assume $|t_l| \geq |t_r|$. In the following, we discuss different cases of whether t_l and t_r are complex or simplex trees. We will show that, in all cases, joining $|t_l|$ and $|t_r|$ has work $O(\log |D| + B)$.

1. When both t_l and t_r are simplex trees. From Theorem 6.1, join costs $O(B)$ work.
2. When both t_l and t_r are complex trees. From Theorem 6.1, join costs $O(\log \frac{|r_l|}{|t_r|}) = O(1)$ work.
3. When t_l is a complex tree, but t_r is a simplex tree. This means that $|t_l| > B$ and $|t_r| \leq B$. From Theorem 6.1, join costs $O\left(B + \log \frac{|t_l|}{B}\right)$ work. Note that, as stated above, $|t_r| \geq \alpha|P|$. Considering $|t_r| \leq B$, we know that $|P| = O(B)$, which also indicates $|t_l| = O(B)$. Plug this into the work of join $O\left(B + \log \frac{|t_l|}{B}\right)$, we can get the bound $O(B)$.

In summary, in all cases, the work of join or join2 is asymptotically bounded by the corresponding split function. \square

From the above two lemmas, we have that the total work of `op` is asymptotically bounded by the split work and unfold work. Next, we prove the bounds for split work and unfold work, respectively.

Lemma 6.10. *The unfold work is $O(\min(mB, n))$.*

Proof. First, note that in our unfolded version of base cases, any block needs to be unfold at most once. Each time the algorithm hits a flat nodes, it unfold the entire subtree in $O(B)$ time.

If $|T_p| = m$, we note that there are at most m/B flat nodes in T_p that needs to be unfolded, so the total work to unfold T_p is $O(m)$. Each of the $O(m)$ entries in T_p will cause an unfold on at most one block in T_d . Therefore, the total work to unfold T_d is $O(mB)$. On the other hand, note that each

block in T_d can be unfolded at most once, which also means that the work of unfolding T_d is $O(n)$. In summary, the work of all unfold functions is $O(\min(mB, n))$.

If $|T_d| = m$ and $|T_p| = n$, there will be at most $O(m)$ nodes in T_p used to process a subset tree in T_d . Since $|T_d| = m$, the total work of unfolding T_d is at most $O(m)$. Based on the same argument as above, the total work of unfolding T_p is $O(mB)$ because at most $O(m)$ unfold functions are invoked, and is also $O(n)$ because there are at most $O(n)$ entries in T_p . The total work is also $O(\min(mB, n))$. \square

Lemma 6.11 (Split work on expanded trees). *The total split work done on two expanded weight-balanced trees of sizes n and $m \leq n$ is $O(m \log \frac{n}{m})$.*

This directly follows [11].

Lemma 6.12. *The total split work is $O(m \log \frac{n}{m} + \min(mB, n))$.*

Proof. The total work for split functions can be viewed as two parts: the total work to done by split functions to traverse and split non-flat nodes, and the work to expose and split the flat nodes. Note that here “non-flat nodes” include both regular nodes in complex trees, and all the nodes in expanded trees.

We first note that in the base cases, the $\text{split}(\cdot, \cdot, \text{true})$ function must be working on a expanded tree. As a result, the total work to traverse and split all non-flat nodes can be asymptotically bounded by the split work considering if both T_p and T_d are fully expanded. This cost can be computed by Lemma 6.11, which is $O(m \log \frac{n}{m})$.

We then consider all work done by the split functions on flat nodes in the non-base cases. We will show it is $O(\min(mB, n))$. Note that in each split, this happens at most once, costing $O(B)$ work. If $|T_p| = m$, there can be $O(m/B)$ such split function calls in the non-base cases, and thus the total non-base case split work on flat nodes is $O(m)$.

If $|T_p| = n$, we discuss in two cases. If $n/B \leq m$, there are $O(n/B)$ regular nodes in T_p , and thus there can be at most $O(n/B)$ split calls. Therefore the total work in this case is $O(n)$, which is also $O(\min(mB, n))$. If $n/B > m$, there are at most $O(m)$ such split function calls, since there are only m nodes in T_d . In this case, the total work of this part is $O(mB)$, which is also $O(\min(mB, n))$. \square

6.2 Non-parallelizable Encoding Schemes

As mentioned at the beginning of this section, if the encoding scheme is not parallelizable, when we deal with a flat node, we have to deal with it sequentially and this can affect the span bound of our algorithms. Again we will use the set algorithms as examples.

Theorem 6.13. *Consider the union algorithm (and similar intersection and difference algorithms) in Fig. 5 on two PaC-trees of sizes m and $n \geq m$ using encoding scheme C . When C takes $O(B)$ work and span to compress and decompress*

a block of size B , the span for these set algorithms is $O(B + \log n'(B + \log m'/B))$, where n' is the pivot tree size, and m' is the decomposed tree size.

Proof. The algorithms need $O(\log |T_p|/B)$ rounds to reach a flat node, where the flat node will be expanded, taking $O(B)$ span. Note that this $O(B)$ additional span is taken only once for each flat node, and they are all at the leaf level of T_p . As a result, they do not add up and can be charged only once in the span. Then the algorithm keeps recursing until a nil node is reached, which is $O(\log B)$ rounds. The total number of rounds of recursive calls is still $O(\log n')$. In each of the recursive calls, we need $O(B + \log m'/B)$ span to deal with split and join. In total the span is $O(B + \log n'(B + \log m'/B))$. \square

7 Work Tradeoff between Updates and Queries

In PaC-trees, we flatten and compress subtrees with size no more than $2B$ where B is a predefined parameter. This approach has several benefits: it is more space efficient, allows for effective compression, and reduces the memory footprint for updates and queries. The only disadvantage is that an insert or delete now has work $O(B + \log n)$, while the P-tree only uses $O(\log n)$ work. We now show an alternative solution for PaC-trees which are updated in-place such that the amortized work for an update is only $O(\log(n/B))$, at a cost of more work for queries (i.e., $O(B + \log n)$).

The basic idea is to leave each leaf node unsorted. In addition, we can keep a linked list for all entries. As such, for an update, we simply find the corresponding leaf node. For an insertion, we add this entry to the end of the linked list of this leaf tree node. For a deletion, we remove this node from the linked list. For both cases, we update the counter for leaf node size, and split or merge if necessary. This also works for inserting or deleting a batch of entries. The only difference is that it is easier to mark tomb entries for batch deletions and physically delete the entries in the next leaf split or merge, since deleting multiple nodes in a linked list is hard (requiring list ranking), while marking tombs can be done when all entries are kept in an array. Here we assume we can locate an entry using its unique identifier, so we can map the identifier to the position in the linked lists. For instance, if the entries are vertices in a graph, then we can use vertex labels as the identifiers for the entries.

This approach increases the lookup cost to $O(B + \log n)$ since now we need to check all entries in a leaf node in the worst case (the work for `aug_range` query is still $O(B + \log n)$ since we only need to check two leaf nodes).

An additional change is the leaf node size. In the previous algorithm, leaf nodes have sizes between B and $2B$. In this setting, we relax it to be B and $(2 + 3c)B$ for any constant $c > 0$. For instance, if $c = 0.1$, then the leaf node size will be between B and $2.3B$. This change is needed to amortizing the

split and merge work for the updates. If $c = 0$, we can have a tree with all leaf nodes containing B entries. If we delete any entry, the associated leaf nodes will contain $B - 1$ entries, need to merge with a neighbor leaf and end up with having $2B - 1$ nodes. Then if we insert two entries in this leaf, we need to split again with $O(B)$ work. To avoid this, we set a padded region of size cB on both sides of the range—once resized, the new leaf node contains $(1 + c)B$ to $(2 + 2c)B$ entries. As such, we need to remove or insert another cB records to trigger the next resizing, so the amortized work is $O(B/cB) = O(1)$ per update.

Theorem 7.1. *A batch of m insertions or deletions can be processed using $O(m \log(n/B))$ amortized work if the batch is unsorted, or $O(m \log(n/Bm))$ amortized work if the batch is sorted. The span is $O(\log(n/B) \log m + \log B)$ for both cases.*

Proof. We first assume the update batch is unsorted. We use the tree root to partition the batch, which takes $O(m)$ work and $O(\log m)$ span. Then we can recursively and in parallel update the left part of the batch and the left subtree, the right part of the batch and the right subtree. After they both finish, we join the two trees with the tree root. The base case is when the corresponding batch for the subtree is empty, or the tree goes to a leaf node. We terminate for the first case. For the second case, we update the leaf with work proportional to the update array size (concatenation for insertions, marking tombs for deletions). The update may trigger a clean-up for the leaf array (split if the array size is larger than $(2 + 3c)B$, merge if the size is smaller than B , or pack if the tombs occupy over a constant fraction of the array). Once a clean-up is triggered, the work is $O(B)$ and the span is $O(\log B)$, and as explained, the work is constant amortized to each previous update to this leaf node.

Similar to the previous analysis, we can split the work into the split work, the join work, and the base case work. The base case work is constant per update. The split work is $O(m \log(n/B))$ —each tree level will partition the update array and the total cost is $O(m)$ per tree level. Since the tree has size n and $O(\log(n/B))$ levels, the total split work is $O(m \log(n/B))$. The join work is logarithmic in the subtree size, and at most m leaf nodes are modified, so the overall work for join is $O(m \log(n/Bm))$, bounded by split work. Putting all pieces together, the work is $O(m \log(n/B))$ amortized. The span is $O(\log(n/B) \log m + \log B) - O(\log m)$ for split and join for $O(\log(n/B))$ levels, and $O(\log B)$ for the base case.

If the update batch is sorted, the split becomes a binary search with cost logarithmic in the current update array size. In this case, the split work is bounded by the join work, so the total work becomes $O(m \log(n/Bm))$ amortized. The span remains unchanged. \square

This new version of the PaC-tree may of interest when updates are more frequent than queries, or the queries are

more costly. For instance, if the query is “reporting the top- k elements”, where k can be large, the work for this query is $\Omega(k)$. In this case, we can use this alternative version of the PaC-tree to reduce the update cost. Assume k is static throughout the algorithm, we can set $B = k$ for the CPAM tree. Based on Theorem 7.1, each update costs $O(\log(n/k))$ work. For a query, we only need to look into the first leaf array which contains at least k entries and at most $(2 + 3c)k$ entries. Then we can use the classic algorithm [34] to find the k -th element from this array, and pack those no larger than it. Hence, this query takes $O(k)$ work and $O(\log k \log \log k)$ span, and the span can be optimized to $O(\log k)$ using the deterministic sampling technique in [13]. This is much more efficient than directly running the classic algorithm [34] with $O(n)$ work, and better than keeping the entire search tree (e.g., a P-tree) which has $O(\log n)$ work per update.

8 Implementation

In this section, we describe CPAM, our implementation of PaC-trees. CPAM is built in C++, based on the PAM framework [52]. Our implementation of sequence and map primitives are mostly unchanged. Most of the changes are to introduce flat nodes, to handle folding and unfolding in join, to express the recursive functions using the expose primitive, and in some cases to add optimized base cases.

Optimized Base Cases. We first implemented union as in Fig. 5, which recursively calls expose to access the left and right subtrees. Although simple and theoretically efficient, in practice unfolding flat nodes into expanded trees and recursing on these trees requires additional memory allocations, and potentially more cache-misses. We therefore designed a new sequential base-case for union when $|T_1| + |T_2| < \kappa$, where κ is a configurable *base-case granularity*. Our base-case works by writing both T_L and T_R into a pre-allocated array A of size κ and merging them in-place to perform the union. It then constructs a PaC-tree from the result in A . Compared to the original version of union that only uses expose, using the special base-case with $\kappa = 4B$ is 4.4x faster, and using $\kappa = 8B$ is 6.7x faster ($B = 128$). We observed similar improvements for some other commonly-used primitives such as filter, map_reduce, multi_insert, multi_delete, and intersection. We use $\kappa = 8B$ in our experiments. We use a parallel granularity of $4B$, which is the threshold for forking parallel tasks in algorithms such as filter and union.

Persistence and Memory Management. CPAM uses a reference counting garbage collector for memory management. CPAM provides functional ordered maps, and thus by default does not modify the input trees. However, in certain cases an application may wish to modify a tree in-place to save memory, e.g., when updates and queries are separated. Although one could deal with in-place and functional updates separately, this is not attractive. Instead, we designed a simple approach to handle both cases using the same code.

Our approach is to store an additional bit indicating whether the supplied node is visible solely to the current function, or whether the node has some external observer, and should therefore be copied. We refer to these special pointers to tree nodes with an additional bit for visibility as *extra pointers*. When an extra pointer is copied, e.g., an algorithm like `union` wishes to use it as part of the resulting tree, we copy this node if the visibility bit is set or if the node has a reference count more than 1, and otherwise we simply return the same node. Similarly, when we expose an extra pointer pointing to a regular node v , we set the visibility bits on the children either if the v 's visibility bit is set, or if v has a reference count greater than 1. If v was visible only to the caller, as an optimization we return it as an additional result, allowing the caller to potentially reuse this node. Our approach lets us write simple algorithms which modify the tree in-place when possible, and begin copying once it reaches subtree that is visible to other observers.

Compression on Blocks. CPAM makes it easy to apply user-specified encoding schemes. Our data structure is templated over a type representing a block encoding scheme (no encoding by default). To add a new encoding scheme, users provide a structure with methods that calculate the encoded size for a block, encode the elements into a buffer, and decode elements from an encoded buffer. This design allows users to specify encoding schemes based on the underlying data type or application, such as text compression. For example, it is easy to add new types of difference coding, e.g., using γ -coding, which would obtain better space usage at the expense of worse running time [49].

9 Applications

In this section we describe four applications that we implement using CPAM. Our inverted index, and range and interval tree applications are based on the implementations from PAM [52]. Our graph processing application is based on Aspen [25]. We focus on the key features of the applications in the context of PaC-trees here.

Inverted Index. We implement a weighted inverted index, similar to those used in search engines. The inverted index maintains a top-level map from words to document lists ($B = 128$). Each document list is a map from document id to an importance score ($B = 128$). The document lists are augmented to maintain the highest importance score. The inverted index supports standard AND/OR queries over words, returning results by rank, and top- k (based on importance) queries. The document ids are compressed using difference encoding, requiring less than two bytes per document.

2D Range Tree. The two-dimensional range tree is a top-level map from x -coordinate to y -coordinate ($B = 128$). The tree is augmented so that every internal node stores all y -coordinates in its subtree (this is itself a set represented as a PaC-tree with $B = 16$). Updates can add and delete

points, and queries can list of or count the points in a given rectangular range. The range tree supports count queries in $O(\log^2 n)$ time, which can be batched to run in parallel.

Interval Tree. The interval tree maintains intervals over the number line, for example, representing the time of a TCP connection, or the time a user is logged into some service. A stabbing query can report all or any intervals that cross a given point. The intervals are represented as an augmented tree from left-coordinate to right-coordinate with $B = 32$. The augmentation maintains the maximum right-coordinate in the subtree. This allows stabbing queries in time $O(k \log n)$ where k is the number of intervals requested or returned (whichever is less). Intervals can be inserted or deleted in $O(\log n)$ time and can be batched to run in parallel.

Graph Processing. Graphs are represented as a two-level structure similar to the inverted index, with a top-level augmented tree (the *vertex tree*) from vertices to edge lists ($B = 64$). Each edge list is a map from neighbor-id to an edge-weight (or empty when unweighted) called an *edge tree* ($B = 64$). The augmentation on the vertex tree maintains the total number of edges in the graph. We focus on unweighted graphs in this paper but note that our implementation also supports weights. As with inverted indices, using difference encoding allows us to store an edge using just 2–3 bytes on average including the bytes used for regular nodes.

On top of this representation, we implement graph algorithms using the Ligna interface [48], including breadth-first search, maximal independent set, and single-source betweenness centrality. Our implementations are based on the ones in Aspen and GBBS [26, 27]. We design parallel batch-updates for our representation, which are applicable in graph-streaming and batch-dynamic graph algorithms.

10 Experiments

Experimental Setup. We run experiments on a 72-core Dell PowerEdge R930 (with two-way hyper-threading) with 4×2.4 GHz Intel 18-core E7-8867 v4 Xeon processors (with a 4800MHz bus and 45MB L3 cache) and 1TB of main memory. Our programs use a work-stealing scheduler for parallelism [10]. We use `numactl -i all` to balance the memory allocations across the sockets for parallel executions. Unless otherwise mentioned, all of the reported numbers are run on 72 cores with hyper-threading.

Overview of Results We show the following experimental results in this section.

- PaC-trees are competitive with PAM for microbenchmarks (Section 10.1) and applications including inverted indices (Section 10.3) and 2D range queries and 1D interval queries (Section 10.4) while using 2.1x–7.8x less space.
- Varying the block size B for an PaC-tree trades off performance for space efficiency (Section 10.1). For even a modest value of $B = 128$, PaC-trees use only 1% more space than a (static) compressed array.

- For graph processing and streaming, CPAM uses 1.3–2.6x less space compared to Aspen, and is almost always faster than Aspen in all tested graph algorithms (Section 10.5).

10.1 PaC-Tree Performance

We begin by studying the performance and space of PaC-trees on a set of microbenchmarks and compare with P-trees from PAM. All experiments in this section use maps and augmented maps where the keys and values are both 64-bit integers. Unless otherwise mentioned PaC-trees use $B = 128$.

Microbenchmark Performance. Table 2 shows the results on PaC-trees, PaC-trees with difference-encoding (DE), and P-trees for a representative subset of the map and sequence primitives. The speedups for both types of PaC-trees range from 28.7–101x and are largest for the version using DE due to additional work for difference encoding. In absolute running time, PaC-trees with DE are usually slower than PaC-trees due to compression and decompression costs, but the overhead is mostly within 10%.

In most of the primitives tested, PaC-trees are faster than P-trees while also using 2.5x less space. For example, PaC-trees are 1.68x faster than P-trees in union on two trees of sizes 10^8 . We note that in this case, the union processes the entirety of both input trees, and so the more cache-friendly processing of blocks in PaC-trees results in lower time. However, if sizes of the two trees are different, the work for union only depends on the smaller size. In this case, since the cost of union using PaC-trees has an additional $O(mB)$ term compared with P-trees, PaC-trees are 5.5x slower than P-trees. However, we expect better performance for smaller block sizes ($B < 128$), which we discuss next.

Effect of Varying B on Performance. Fig. 12 shows the results of varying the block size B , on the performance of various operations. Most operations obtain speedups as B is increased up until $B = 16$. For the sequential operations, such as find and range, we see a steady increase in the running time for $B > 16$ and see a similar trend for Union-Imbal, which takes the union of trees with 10^8 and 10^5 elements. This slowdown with increasing B is due to the extra $O(mB)$ term in the work of union. For the smallest block size ($B = 1$), our running time matches that of P-trees on this operation.

Space Usage. For $B = 128$, PaC-trees obtain a 2.48x reduction in space usage compared to using P-trees, and a further 1.73x reduction in space usage by using difference encoding. The 10^8 pairs stored in the experiments require 1.6GB of memory to represent as a single flat array, which is also a lower bound for the space usage of a search tree structure. To understand how close PaC-trees come to this lower bound, we study the space usage of unaugmented maps using PaC-trees as a function of the block size B (Fig. 13). Using $B = 32$, PaC-trees are only 1.05x larger than the lower bound and using $B = 128$, it is just 1.01x larger than the lower bound. For $B = 128$, just 1.1% of the allocated memory is used for

regular nodes and metadata in the flat nodes. These savings are obtained without using any additional encoding. Applying difference encoding improves the space by 1.77x over the unencoded trees and the array lower bound, and is only 1.03x larger than the space used to difference encode all of the keys in a single array, leaving the values uncompressed, which is a lower bound for a search tree structure using difference encoding for such input.

Using PaC-trees requires much lower space overhead for augmentation compared to P-trees (Fig. 13). For P-trees, adding 8 byte augmented values increases the size of the maps by 20%, whereas PaC-trees (both with and without difference encoding) using $B = 128$ incurs only a 1% increase in space for the augmented values. The savings comes from only storing a single augmented value per flat node, which only uses extra space proportional to n/B augmented values.

10.2 Comparison with Collections in Spark

We compared CPAM with the shared-memory parallel implementation of Apache Spark on a simple benchmark drawn from the Apache Spark tutorial. The benchmark first loads the same Wikipedia corpus that we use for our inverted index application (1,943,575,146 words in 8,125,326 documents). The first example then tokenizes the dataset into words, and computes the longest word length. The second example computes the most frequently occurring word by using the reduceByKey primitive in Spark to group common words and compute the mostly frequently occurring word using a reduce. We use in-memory caching for the intermediate mapped dataset and report the fastest (cached) time.

For the first example, Spark takes 46.9 seconds for the first (uncached) run, and the subsequent average (cached) time is 21.5 seconds. Our CPAM implementation, where the dominant cost is the memory-bound parsing step, requires 6.57 seconds on average (3.2x faster than the cached time). On the second example, Spark takes 96.3 seconds for the first run, and 72.5 seconds for the average cached time. For the second example, the dominant cost for CPAM is the parallel sort (we use a parallel sample-sort). The end-to-end time is 14.6 seconds, which is 4.9x faster than the Spark cached time.

We also tried evaluating the same set of sequence benchmarks shown in Fig. 2 using Spark, but observed significantly worse running times for all of the sequence primitives (up to 2 orders of magnitude worse performance; e.g., reduce on a 100M element sequence with 8-byte elements takes 2.07 seconds, whereas CPAM takes 0.00865 seconds). The slowdown could be due to fixed parallelization overheads in Spark, although their word counting example which we studied above performs reasonably well despite working over a significantly larger dataset (nearly 2 billion words).

10.3 Inverted Index

Next, we study our performance on the inverted index application. We run the application on documents derived from a

	n	m	PaC-tree			PaC-tree (Diff)			P-tree (PAM)		
			T_1	T_{144}	Spd.	T_1	T_{144}	Spd.	T_1	T_{144}	Spd.
No augmentation											
Size (GB)	10^8	—	1.61	—	—	0.926	—	—	4.00	—	—
Build	10^8	—	5.55	0.186	29.8	5.71	0.180	31.7	5.94	0.221	26.8
Union	10^8	10^8	5.33	0.088	60.5	6.29	0.089	70.6	8.97	0.168	53.3
Union	10^8	10^5	1.09	0.021	51.9	1.28	0.022	58.1	0.206	0.0038	54.2
Intersect	10^8	10^8	4.35	0.065	66.9	5.68	0.081	70.1	9.50	0.139	68.3
Difference	10^8	10^8	3.00	0.055	54.4	3.55	0.056	63.3	8.17	0.123	66.4
Map	10^8	10^8	0.859	0.037	22.9	1.14	0.023	49.5	1.32	0.091	14.5
Reduce	10^8	—	0.306	0.018	17.0	0.308	0.0092	33.4	1.60	0.034	47.0
Filter	10^8	—	0.997	0.028	35.6	1.24	0.018	68.8	1.90	0.0524	36.2
Find	10^8	10^8	103	1.17	88.0	125	1.23	101.6	105.5	1.05	100.4
Insert	10^8	10^6	0.829	—	—	1.42	—	—	0.773	—	—
Multi-Insert	10^8	10^8	18.8	0.332	56.6	19.9	0.323	61.6	9.67	0.338	28.6
Range	10^8	10^6	11.5	0.318	36.1	13.1	0.226	57.9	3.77	0.0738	45.6
With augmentation											
Size (GB)	10^8	—	1.63	—	—	0.936	—	—	4.80	—	—
Build	10^8	—	5.66	0.197	28.7	5.84	0.186	31.3	6.48	0.246	26.3
Union	10^8	10^8	5.52	0.098	56.3	6.52	0.090	72.4	10.13	0.196	51.6
AugRange	10^8	10^7	12.3	0.331	37.1	13.9	0.234	59.4	4.80	0.082	58.5
AugFilter	10^8	—	12.2	0.333	36.6	13.6	0.234	58.1	4.95	0.081	61.1

Table 2. Microbenchmark results. We fix $B = 128$ for PaC-trees. n is the tree size. For set functions and multi-insert, $m \leq n$ is the size of the other set (batch). For other functions, m is the number of queries tested. T_1 is the sequential running time. T_{144} is parallel running time using 72 cores (144 hyperthreads). *Diff* means difference encoding. We highlight the best parallel running time (or size) per experiment in green and underlined.

large Wikipedia dataset also used by PAM for a fair comparison. The dataset is processed by removing all markup, converting characters that are not alphanumeric to whitespace and making all words case insensitive [52]. The processed dataset contains 1.94 billion words over 8.13 million documents. Like PAM, our evaluation measures the performance of (1) building an index over (words, doc_id, weight) triples and (2) running queries that fetch the posting lists for two words, compute the intersection of the lists, and select the top 10 documents by weight.

Table 3 shows the results of the experiment. For building the index, our implementation achieves 76x speedup and our parallel running times are comparable with those of PAM (at most 1.1x slower). For the queries, we observe that the unencoded trees achieve essentially the same parallel time as PAM, whereas the difference encoded trees are 1.18x slower due to the higher cost of intersection operations in our difference encoded implementation. The space usage using PaC-trees is much smaller than that of PAM, being 3.84x smaller without encoding and 7.81x smaller using a custom encoder that combines difference encoding for the keys with byte-encoding for the integer values (weights).

10.4 Interval and Two-Dimensional Range Trees

We benchmark our interval and two-dimensional range trees as in PAM [50]. We build our interval tree on 10^8 intervals, and for queries run stabbing queries over 10^8 points

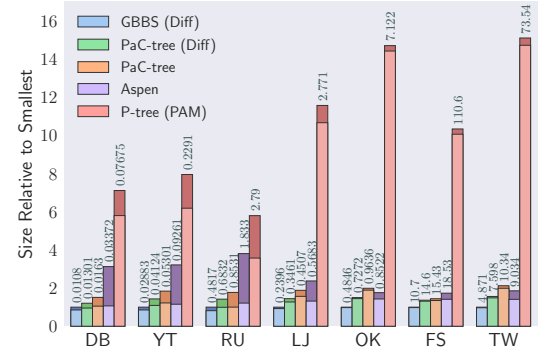


Figure 11. Relative space usage of different graph representations. GBBS (Diff) is our *static* baseline compressed graph representation. PaC-tree uses PaC-trees for vertex and edge trees, and PaC-tree (Diff) difference encodes both trees. Aspen uses P-trees for the vertex tree and C-trees with difference encoding for edge trees. P-tree (PAM) uses P-trees for the vertex and edge trees. The values on top of each bar are the memory usage in GiB.

	Library	Space	Method	n	m	T_1	T_{144}	Spd.
Inverted Index	PaC-tree	8.29	Build	10^8	—	746	9.73	76.6
			Query	10^8	10^8	341	<u>4.46</u>	76.4
	PaC-tree (D)	<u>4.07</u>	Build	10^8	—	754	9.81	76.8
			Query	10^8	10^8	367	5.32	68.9
Interval	P-tree (PAM)	31.9	Build	10^8	—	575	<u>8.86</u>	64.9
			Query	10^8	10^8	313	4.48	69.8
	PaC-tree	<u>0.812</u>	Build	10^8	—	10.9	<u>0.179</u>	60.8
			Query	10^8	10^8	60.8	<u>0.525</u>	115.8
Range	P-tree (PAM)	3.54	Build	10^8	—	11.6	0.271	42.8
			Query	10^8	10^8	54.3	0.628	86.4
	PaC-tree	<u>40.3</u>	Build	10^8	—	164	<u>2.71</u>	60.7
			Q-Sum	10^8	10^6	54.2	<u>0.629</u>	86.1
			Q-All	10^8	10^3	7.20	<u>0.266</u>	27.0
	P-tree (PAM)	89.6	Build	10^8	—	169	2.84	59.6
			Q-Sum	10^8	10^6	60.7	0.735	82.5
			Q-All	10^8	10^3	21.6	0.552	39.1

Table 3. Build and query times and space usage in GiB for inverted index, interval tree, and range tree applications. T_1 is the single-thread time, T_{144} is the 72-core time using hyper-threading, and Spd. is the parallel speedup. The best parallel running time (or size) is highlighted in green and underlined per experiment.

in parallel. We observe that both building and querying the trees achieves good parallel speedup (60–115x). PaC-trees are 1.51x faster than PAM in construction, and is 1.19x faster

Graph	Vertices	Edges	Ours	Aspen	Aspen Ours
<i>DBLP (DB)</i>	425,957	2,099,732	0.0130	0.03409	2.62x
<i>YouTube (YT)</i>	1,138,499	5,980,886	0.0412	0.0934	2.26x
<i>USA-Road (RU)</i>	23,947,348	57,708,624	0.683	1.843	2.69x
<i>LiveJournal (LJ)</i>	4,847,571	85,702,474	0.346	0.527	1.52x
<i>com-Orkut (CO)</i>	3,072,627	234,370,166	0.727	0.893	1.22x
<i>Twitter (TW)</i>	41,652,231	2,405,026,092	7.59	9.42	1.23x
<i>Friendster (FS)</i>	65,608,366	3,612,134,270	14.6	19.1	1.30x

Table 4. Statistics about tested graphs and memory usage of PaC-tree and Aspen in GiB.

for queries. Overall we find that PaC-trees enable better performance than PAM while using 4.37x less space.

We build our range trees on 10^8 uniformly random points in the plane between $(0, 0)$ and $(1e8, 1e8)$. We run two types of queries: the first count the number of points in the range (Q-Sum), and the second returns all points in the range. We tuned the window sizes used in our queries to match the settings evaluated by PAM (around 10^6 points returned per query). Both PaC-trees and P-trees build the data structure in a similar amount of time. PaC-trees achieve better performance than P-trees for both queries, being 1.16x faster for Q-Sum and 1.96x faster for Q-All queries, likely due to requiring fewer cache-misses when processing the tree to output the points within a given range. The range tree application using PAM has previously been compared with range trees in CGAL [45] and was shown to outperform it [50].

For space usage, PaC-trees result in 2.18x less space compared to PAM. We note that 95% of the space used in PAM is for the P-trees stored as augmented values in each node (representing the union of the y -coordinates in the subtree). The majority of our savings come from compressing the augmented trees using PaC-trees which results in a 2.53x less space for the inner trees, and 2.18x less space overall.

10.5 Graph Processing and Graph Streaming

Our last set of experiments study the performance of PaC-trees for a set of standard benchmarks from the graph processing and graph streaming literature. Our evaluation roughly follows Aspen’s and we compare our performance and space usage with that of Aspen and its C-tree implementation.

Graph Data. *DBLP* is co-authorship network based on research papers in computer science. *YouTube* is a social network graph based on YouTube. *USA-Road (RO)* is an undirected road network from the DIMACS challenge [24]. *LiveJournal (LJ)* is a directed graph of the LiveJournal social network [17]. *com-Orkut (CO)* is an undirected graph of the Orkut social network. *Twitter (TW)* is a directed graph of the Twitter network [40]. *Friendster (FR)* is an undirected graph describing friendships from a gaming network. The DBLP, YouTube, and Friendster graphs are obtained from the SNAP dataset [41]. We note that some of our inputs (like the LiveJournal graph) are originally directed, and we symmetrize them before applying our algorithms to maintain

Graph	Aspen		Ours				Aspen Ours
	FS	FS Time	No-FS	FS	FS No-FS	FS Time	
BFS	<i>LiveJournal</i>	21.7	3.82	19.8	17.5	1.13x	1.38
	<i>com-Orkut</i>	15.3	2.35	14.5	12.4	1.16x	1.12
	<i>Twitter</i>	138	37.8	125	112	1.11x	1.25
MIS	<i>LiveJournal</i>	55.3	3.82	72.0	45.7	1.57x	1.38
	<i>com-Orkut</i>	70.2	2.35	96.9	69.2	1.40x	1.12
	<i>Twitter</i>	1022	37.8	1190	971	1.22x	12.5
BC	<i>LiveJournal</i>	74.6	3.82	82.1	72.3	1.13x	1.38
	<i>com-Orkut</i>	76.3	2.35	88.6	78.2	1.13x	1.12
	<i>Twitter</i>	1150	37.8	2735	1030	2.65x	12.5

Table 5. Parallel running times (in milliseconds) for Aspen and our implementation. We show the algorithm performance without flat snapshots (**No-FS**), with flat snapshots (**FS**), and the time to computing the flat snapshot (**FS Time**).

consistency with prior work on Aspen [25] and GBBS [26] that symmetrize graphs in their evaluations. Table 4 shows information about our graph inputs, including the number of vertices, edges, and space used.

We evaluate five graph representations including using PAM, Aspen, PaC-tree with or without difference encoding, and GBBS. Aspen uses C-trees as edge trees and leaves vertex trees uncompressed using P-trees. GBBS is a state-of-the-art static graph processing library which represents graphs as static arrays using difference encoding, which serves as our baseline of graph representation. Fig. 11 shows the relative size of each graph format. We see that the smallest format in all cases is **PaC-tree (Diff)**, which applies PaC-trees with difference encoding for both vertex and edge trees. Using this format yields a space improvement of between 4–9.7x over just using P-trees. For the graphs with high average-degree, most of the savings come from using PaC-trees for the edge trees. Adding difference encoding to both trees yields between 1.05–1.32x space improvement. PaC-trees are also 1.3–2.6x more space-efficient than Aspen. Note that C-trees in Aspen are also difference encoded, so the main difference between the two representations is that PaC-tree (Diff) also uses PaC-trees to chunk the vertex tree, and that PaC-trees employ a deterministic strategy for chunking. PaC-trees with difference encoding achieves consistently lower space compared with Aspen, ranging between 1.3x for Friendster, our largest graph, to a maximum space improvement for 2.62x on USA-Road, our sparsest graph. The space savings come from chunking the vertex trees, which is not possible in Aspen, since the C-tree implementation is specialized for edge trees.

Graph Algorithm Performance. We study the performance of three fundamental graph kernels: breadth-first search (BFS), single-source betweenness centrality (BC), and maximal independent set (MIS). Our implementations are based on those in Aspen. We study performance using our most space-efficient version (PaC-tree (Diff)). Following Aspen, our implementation also supports the *flat snapshot* object, which is an array storing all vertices in the current graph. The

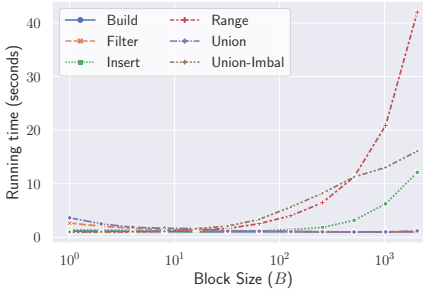


Figure 12. Primitive running times for PaC-trees vs. block size B . We use 10^8 key-value pairs (8 bytes each). Union, Intersection and Difference all work on two trees with 10^8 elements. Union-Imbal takes the union of trees with 10^8 and 10^5 elements.

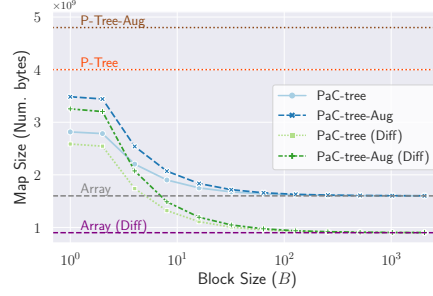


Figure 13. Size of PaC-trees (with or without DE) as a function of block size B . We use 10^8 key-value pairs (8 bytes each). For augmented maps (-Aug), augmented values are 8 bytes each. The grey line shows the number of bytes to store the 10^8 elements in an array and the purple line shows the bytes used to store the difference encoded keys in a single array using byte encoding.

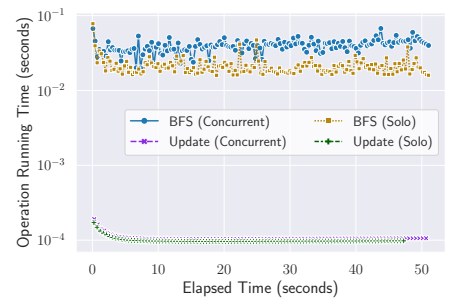


Figure 14. Performance of concurrent updates and queries. The time series plot illustrates running times when running BFS queries with batch-insertions of edges concurrently (Concurrent), and when queries and updates are run individually (Solo) on the LiveJournal graph.

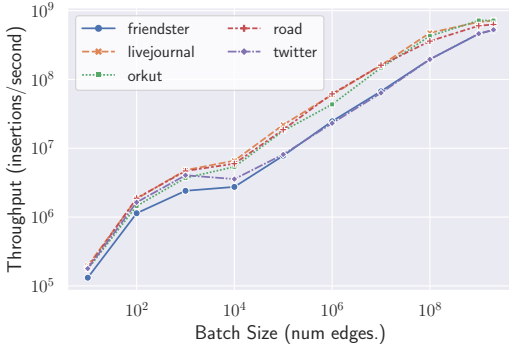


Figure 15. Edge insertion throughput (insertions per second) for our graph update algorithms as a function of the number of edge insertions in each batch. The throughput for batch deletions is similar to that of insertions.

idea is that instead of accessing edges for a vertex through the vertex tree (performing tree traversal), algorithms directly access edge trees through the flat snapshot.

Table 5 shows performance results for three of our graph datasets. Across all three kernels our implementations are 1.12x faster than Aspen’s implementations on average. We observe that flat snapshots can be generated 2.09–3.02x faster in CPAM due to PaC-trees requiring fewer cache-misses to traverse than P-trees when creating flat snapshot array. We note that the implementation of edgeMap and other primitives from Ligra (including constants and other tuning parameters) are exactly the same in both CPAM and Aspen. Aspen also difference encodes in its edge trees (represented using C-trees). The performance improvements that we observe are therefore a result of PaC-trees providing faster flat snapshots, and having better balance in chunk sizes compared to the randomized approach used in C-trees.

Graph Batch-Update Throughput In this section we study the performance of our graph representation using PaC-trees when performing updates in varying size batches. We focus on algorithms for inserting and deleting batches of edges, since inserting vertices (along with incident edges) can easily be done using the edge insert primitive, and vertex deletions simply use `multi_delete` on PaC-trees. Our experiment follows the methodology used in Aspen. To generate updates, we sample directed edges from an rMAT generator [21] with $a = 0.5, b = c = 0.1, d = 0.3$. We use the same generator for our concurrent update and query experiment in Section 10.5. For a batch of size K , we generate K directed edge updates from the stream (note that there can be duplicates) and repeatedly insert the edges and delete the edges from the batch, reporting the median of three trials. We note that the times we report include the time to sort and remove duplicates from the batch.

Fig. 15 shows the throughput of batch edge insertions (insertions per second) as a function of the batch size. We note that the throughput for deletions are close to that of insertions (within 10% across all graphs). To remove clutter, we show results on the five largest graphs in our datasets. We observe that the throughput of our graph representation improves with increasing batch size; for the largest batch size, the algorithm achieves a maximum throughput of between 719M edge insertions per second for the com-Orkut graph, and a minimum of 527M edge insertions per second for the Twitter graph. We compared these results with those of Aspen on the same machine and find that we obtain 1.62x higher throughput across the three graphs both systems consider in this experiment, and an average throughput increase of 1.65x across these graphs.

Concurrent Updates and Queries. Our last experiment concurrent updates and queries on graphs. The experiment

performs n undirected edge insertions drawn from the rMAT generator described above. We use a batch size of 5 in the updates (10 directed edges are inserted per batch). We then spawn two parallel jobs, one performing the updates one batch after the other, and the other performing BFS queries, one after the other. Both the updates and queries are parallel (i.e., they internally make use of parallelism).

Fig. 14 shows the result of the experiment. We find that the concurrent queries are 1.85x slower on average than the queries in isolation, and that the concurrent updates are 1.07x slower on average than updates in isolation. In the concurrent setting, the average latency to make one of the update batches visible is 100 microseconds, and the updates achieve a throughput of 94,000 undirected edge updates per second. We leave further optimizations and a more in depth study of the graph setting for future work with our system.

11 Conclusion

We have presented PaC-tree, a deterministic compressed ordered map data structure and an implementation of the structure in a library CPAM. The important features of PaC-trees and its implementation in CPAM include the following.

- It is purely functional allowing for persistent snapshots while updates are being made, and safe for parallelism.
- It supports sequences, ordered sets, ordered maps, and augmented maps, with a wide variety of functions on them.
- It provides theoretical bounds on work, span, and space.
- It achieves fast sequential time and gets up to 100x speedup on 72 cores with 144 hyperthreads.
- It achieves memory usage that is close to a compressed array and up to an order of magnitude smaller than PAM.
- It is internally memory managed using reference counting.
- It is backward compatible with PAM.
- It has been used to implement the full functionality of Aspen while improving runtime and/or space.

For future work, we are interested in extending PaC-trees to support higher-fanout internal nodes, similar to B -trees, which would allow users to improve query latency at the expense of increased work when performing updates. Other future work includes applying PaC-trees to improve space utilization in databases, and to improve the performance of collection-based applications using non-volatile memory.

Acknowledgement

This work was supported by the National Science Foundation grants CCF-1901381, CCF-1910030, CCF-1919223, CCF-2103483, and CCF-2119352.

References

- [1] Umüt A. Acar, Arthur Charguéraud, and Mike Rainey. 2014. Theory and Practice of Chunked Sequences. In *European Symposium on Algorithms (ESA)*.
- [2] Stephen Adams. 1993. Efficient sets—a balancing act. *Journal of functional programming* 3, 04 (1993).
- [3] Vitaly Aksenov, Vincent Gramoli, Petr Kuznetsov, Anna Malova, and Srivatsan Ravi. 2017. A concurrency-optimal binary search tree. In *European Conference on Parallel Processing (Euro-Par)*. Springer.
- [4] Maya Arbel-Raviv, Trevor Brown, and Adam Morrison. 2018. Getting to the Root of Concurrent Binary Search Tree Performance. In *USENIX Annual Technical Conference*.
- [5] Dmitry Basin, Edward Bortnikov, Anastasia Braginsky, Guy Golan-Gueta, Eshcar Hillel, Idit Keidar, and Moshe Sulamy. 2017. KiWi: A key-value map for scalable real-time analytics. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
- [6] R. Bayer and E. M. McCreight. 1972. Organization and maintenance of large ordered indexes. *Acta Informatica* 1, 3 (01 Sep 1972).
- [7] Naama Ben-David, Guy E. Blelloch, Panagioti Fatourou, Eric Ruppert, Yihan Sun, and Yuanhao Wei. 2021. Space and Time Bounded Multi-version Garbage Collection. In *International Symposium on Distributed Computing (DISC)*. <https://doi.org/10.4230/LIPIcs.DISC.2021.12>
- [8] Jean-Philippe Bernardy. 2008. The Haskell Yi package. <http://hackage.haskell.org/package/yi-0.6.2.3/docs/src/Data-Rope.html>.
- [9] Daniel K. Blandford and Guy E. Blelloch. 2004. Compact Representations of Ordered Sets. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [10] Guy E. Blelloch, Daniel Anderson, and Laxman Dhulipala. 2020. ParlayLib - A Toolkit for Parallel Algorithms on Shared-Memory Multi-core Machines. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [11] Guy E. Blelloch, Daniel Ferizovic, and Yihan Sun. 2016. Just Join for Parallel Ordered Sets. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [12] Guy E. Blelloch, Jeremy T. Fineman, Yan Gu, and Yihan Sun. 2020. Optimal Parallel Algorithms in the Binary-Forking Model. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [13] Guy E. Blelloch, Phillip B. Gibbons, and Harsha Vardhan Simhadri. 2010. Low Depth Cache-Oblivious Algorithms. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [14] Guy E. Blelloch, Yan Gu, Julian Shun, and Yihan Sun. 2018. Parallel Write-Efficient Algorithms and Data Structures for Computational Geometry. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [15] Robert D. Blumofe and Charles E. Leiserson. 1999. Scheduling multi-threaded computations by work stealing. *J. ACM* 46, 5 (1999).
- [16] Hans-J. Boehm, Russ Atkinson, and Michael Plass. 1995. Ropes: An Alternative to Strings. *Softw. Pract. Exper.* 25, 12 (1995).
- [17] Paolo Boldi and Sebastiano Vigna. 2004. The WebGraph Framework I: Compression Techniques. In *International World Wide Web Conference (WWW)*.
- [18] Anastasia Braginsky and Erez Petrank. 2012. A lock-free B+ tree. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
- [19] Richard P. Brent. 1974. The Parallel Evaluation of General Arithmetic Expressions. *J. ACM* 21, 2 (April 1974), 201–206.
- [20] Nathan G. Bronson, Jared Casper, Hassan Chafi, and Kunle Olukotun. 2010. A Practical Concurrent Binary Search Tree. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
- [21] Deepayan Chakrabarti, Yiping Zhan, and Christos Faloutsos. 2004. R-MAT: A Recursive Model for Graph Mining. In *SIAM International Conference on Data Mining (SDM)*.
- [22] Raymond Cheng, Ji Hong, Aapo Kyröla, Youshan Miao, Xuetian Weng, Ming Wu, Fan Yang, Lidong Zhou, Feng Zhao, and Enhong Chen. 2012. Kineograph: taking the pulse of a fast-changing and connected world. In *ACM European Conference on Computer Systems (EuroSys)*.
- [23] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms (3rd edition)*. MIT Press.
- [24] Camil Demetrescu, Andrew Goldberg, and David Johnson. 2019. 9th DIMACS Implementation Challenge: Shortest Paths.

- <http://www.dis.uniroma1.it/challenge9/>.
- [25] Laxman Dhulipala, Guy E Blelloch, and Julian Shun. 2019. Low-latency graph streaming using compressed purely-functional trees. In *ACM Conference on Programming Language Design and Implementation (PLDI)*.
 - [26] Laxman Dhulipala, Guy E. Blelloch, and Julian Shun. 2021. Theoretically Efficient Parallel Graph Algorithms Can Be Fast and Scalable. *ACM Transactions on Parallel Computing (TOPC)* 8, 1 (2021). <https://doi.org/10.1145/3434393>
 - [27] Laxman Dhulipala, Jessica Shi, Tom Tseng, Guy E. Blelloch, and Julian Shun. 2020. The Graph Based Benchmark Suite (GBBS). In *Intl. Workshop on Graph Data Management Experiences and Systems (GRADES)*.
 - [28] Faith Ellen, Panagiota Fatourou, Eric Ruppert, and Franck van Breugel. 2010. Non-blocking binary search trees. In *ACM Symposium on Principles of Distributed Computing (PODC)*.
 - [29] Panagiota Fatourou, Elias Papavasileiou, and Eric Ruppert. 2019. Persistent non-blocking binary search trees supporting wait-free range queries. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*.
 - [30] Matthew Fluet, Mike Rainey, John Reppy, and Adam Shaw. 2008. Implicitly-threaded Parallelism in Manticore. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*.
 - [31] Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. 1998. The implementation of the Cilk-5 multithreaded language. *ACM Conference on Programming Language Design and Implementation (PLDI)*.
 - [32] Yan Gu, Yihan Sun, and Guy E. Blelloch. 2018. Algorithmic Building Blocks for Asymmetric Memories. In *European Symposium on Algorithms (ESA)*.
 - [33] Switzerland International Organization for Standardization, Geneva. 2018. ISO/IEC TS 19570:2018: Programming Languages – Technical Specification for C++ Extensions for Parallelism. <https://www.iso.org/standard/70588.html>.
 - [34] J. JaJa. 1992. *Introduction to Parallel Algorithms*. Addison-Wesley Professional.
 - [35] Java Fork-Join, Oracle Java Documentation [n.d.]. <http://docs.oracle.com/javase/tutorial/essential/concurrency/forkjoin.html>.
 - [36] Haim Kaplan and Robert Endre Tarjan. 1996. Purely Functional Representations of Catenable Sorted Lists. In *ACM Symposium on Theory of Computing (STOC)*.
 - [37] Alfons Kemper, Thomas Neumann, Jan Finis, Florian Funke, Viktor Leis, Henrik Mühe, Tobias Mühlbauer, and Wolf Rödiger. 2013. Processing in the Hybrid OLTP & OLAP Main-Memory Database System HyPer. *IEEE Data Eng. Bull.* 36, 2 (2013).
 - [38] Edward A. Kmett. 2010. The Haskell Rope package.
 - [39] H. T. Kung and Philip L. Lehman. 1980. Concurrent Manipulation of Binary Search Trees. *ACM Trans. Database Syst.* 5, 3 (1980).
 - [40] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a Social Network or a News Media?. In *International World Wide Web Conference (WWW)*.
 - [41] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
 - [42] Peter Macko, Virendra J Marathe, Daniel W Margo, and Margo I Seltzer. 2015. LLAMA: Efficient graph analytics using large multiversioned arrays. In *IEEE International Conference on Data Engineering (ICDE)*.
 - [43] Colt McAnlis and Aleks Haekey. 2016. *Understanding Compression*. O'Reilly Media, Inc.
 - [44] Aravind Natarajan and Neeraj Mittal. 2014. Fast Concurrent Lock-Free Binary Search Trees. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
 - [45] Gabriele Neyer. 2017. dD Range and Segment Trees. In *CGAL User and Reference Manual* (4.10 ed.). CGAL Editorial Board. <http://doc.cgal.org/4.10/Manual/packages.html>
 - [46] Jürg Nievergelt and Edward M Reingold. 1973. Binary search trees of bounded balance. *SIAM J. on Computing* 2, 1 (1973).
 - [47] Chris Okasaki. 1999. *Purely functional data structures*. Cambridge University Press.
 - [48] Julian Shun and Guy E. Blelloch. 2013. Ligra: A Lightweight Graph Processing Framework for Shared Memory. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
 - [49] Julian Shun, Laxman Dhulipala, and Guy E Blelloch. 2015. Smaller and faster: Parallel processing of compressed graphs with Ligra+. In *IEEE Data Compression Conference (DCC)*.
 - [50] Yihan Sun and Guy E Blelloch. 2019. Parallel Range, Segment and Rectangle Queries with Augmented Maps. In *Algorithm Engineering and Experiments (ALENEX)*.
 - [51] Yihan Sun, Guy E Blelloch, Wan Shen Lim, and Andrew Pavlo. 2019. On supporting efficient snapshot isolation for hybrid workloads with multi-versioned indexes. *Proceedings of the VLDB Endowment (PVLDB)* 13, 2 (2019).
 - [52] Yihan Sun, Daniel Ferizovic, and Guy E Blelloch. 2018. PAM: Parallel Augmented Maps. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
 - [53] Yuanhao Wei, Naama Ben-David, Guy E. Blelloch, Panagiota Fatourou, Eric Ruppert, and Yihan Sun. 2021. Constant-time snapshots with applications to concurrent data structures. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
 - [54] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J. Franklin, Ali Ghodsi, Joseph Gonzalez, Scott Shenker, and Ion Stoica. 2016. Apache Spark: a unified engine for big data processing. *Commun. ACM* 59, 11 (2016).