

Parallel Combining: Benefits of Explicit Synchronization

Vitaly Aksenov

ITMO University, Saint-Petersburg, Russia and Inria, Paris, France
aksenov@corp.ifmo.ru

Petr Kuznetsov

LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France
petr.kuznetsov@telecom-paristech.fr

Anatoly Shalyto

ITMO University, Saint-Petersburg, Russia
shalyto@mail.ifmo.ru

Abstract

A *parallel batched* data structure is designed to process synchronized *batches* of operations on the data structure using a parallel program. In this paper, we propose *parallel combining*, a technique that implements a *concurrent* data structure from a parallel batched one. The idea is that we explicitly synchronize concurrent operations into batches: one of the processes becomes a *combiner* which collects concurrent requests and initiates a parallel batched algorithm involving the owners (*clients*) of the collected requests. Intuitively, the cost of synchronizing the concurrent calls can be compensated by running the parallel batched algorithm.

We validate the intuition via two applications. First, we use parallel combining to design a concurrent data structure optimized for *read-dominated* workloads, taking a dynamic graph data structure as an example. Second, we use a novel parallel batched *priority queue* to build a concurrent one. In both cases, we obtain performance gains with respect to the state-of-the-art algorithms.

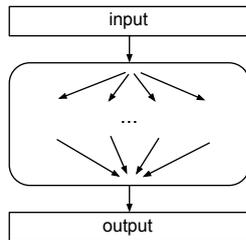
2012 ACM Subject Classification Computing methodologies → Concurrent computing methodologies, Parallel computing methodologies; Theory of computation → Models of computation → Concurrency → Distributed computing models, Parallel computing models

Keywords and phrases concurrent data structure, parallel batched data structure, combining

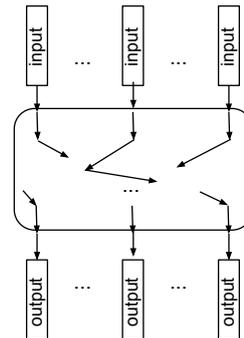
1 Introduction

To ensure correctness of concurrent computations, various synchronization techniques are employed. Informally, synchronization is used to handle conflicts on *shared data*, e.g., resolving data races, or *shared resources*, e.g., allocating and deallocating memory. Intuitively, the more sophisticated conflict patterns a concurrent program is subject to—the higher are the incurred synchronization costs.

Let us consider a concurrent-software class which we call *parallel programs*. Provided an input, a parallel program aims at computing an output that satisfies a *specification*, i.e., an input-output relation (Figure 1). To boost performance, the program distributes the computation across multiple parallel processes. Parallel programs are typically written for two environments: for *static multithreading* and *dynamic multithreading* [11]. In *static multithreading*, each process is given its own program and these programs are written as a composition of *supersteps*. During a superstep, the processes perform conflict-free individual computations and, when done, synchronize to accumulate the results. In *dynamic multithreading*, the program is written using dynamically called *fork-join* mechanisms (or similar ones, e.g., `#pragma omp parallel` in OpenMP [6]). In both settings, synchronization only appears in a specific form: memory allocation/deallocation, and aggregating superstep computations or thread scheduling [17].



■ **Figure 1** Execution of a parallel program



■ **Figure 2** Execution on a concurrent data structure

General-purpose *concurrent data structures*, such as stacks, binary search trees and priority queues, operate in a much less controlled environment. They are programmed to accept and process asynchronous operation calls, which come from multiple concurrent processes and may interleave arbitrarily. If we treat operation calls as inputs and their responses as outputs, we can say that inputs and outputs are distributed across the processes (Figure 2). It is typically expected that the interleaving operations match the high-level sequential semantics of the data type [24], which is hard to implement efficiently given diverse and complicated data-race patterns often observed in this kind of programs. Therefore, designing efficient and correct concurrent data structures requires a lot of ingenuity from the programmer. In particular, one should strive to provide the “just right” amount of synchronization. *Lock-based* data structures obviate data races by using fine-grained locking ensuring that contested data is accessed in a mutually exclusive way. *Wait-free* and *lock-free* data structures allow data races but mitigate their effects by additional mechanisms, such as *helping* where one process may perform some work on behalf of other processes [23].

As parallel programs are written for a restricted environment with simple synchronization

patterns, they are typically easier to design than concurrent data structures. In this paper, we suggest benefiting from this complexity gap by building concurrent data structures *from* parallel programs. We describe a methodology of designing a concurrent data structure from its *parallel batched* counterpart [2]. A parallel batched data structure is a special case of a parallel program that accepts *batches* (sets) of operations on a given sequential data type and executes them in a parallel way. In our approach, we *explicitly* synchronize concurrent operations, assemble them into batches, and apply these batches on an *emulated* parallel batched data structure.

More precisely, concurrent processes share a set of *active requests* using any combining algorithm [31, 21, 15, 16]. One of the processes with an active request becomes a *combiner* and forms a *batch* from the requests in the set. Under the coordination of the combiner, the owners of the collected requests, called *clients*, apply the requests in the batch to the parallel batched data structure. As we show, this technique becomes handy when the overhead of *explicitly synchronizing* operation calls in batches is compensated by the advantages of involving clients into the computation using the parallel batched data structure.

We discuss two applications of parallel combining and experimentally validate performance gains. First, we design concurrent implementations optimized for *read-dominated* workloads given a sequential data structure. Intuitively, updates are performed sequentially and read-only operations are performed by the clients in parallel under the coordination of the combiner. In our performance analysis, we considered a *dynamic graph* data structure [25] that can be accessed for adding and removing edges (updates), as well as for checking the connectivity between pairs of vertices (read-only). Second, we apply parallel combining to *priority queue* that is subject to sequential bottlenecks for minimal-element extractions, while most insertions can be applied concurrently. As a side contribution, we propose a novel parallel batched priority queue, as no existing batched priority queue we are aware of can be efficiently used in our context. Our performance analysis shows that implementations based on parallel combining may outperform state-of-the-art algorithms.

Structure. The rest of the paper is organized as follows. In Section 2, we give preliminary definitions. In Section 3, we outline parallel combining technique. In Sections 4 and 5, we present applications of our technique. In Section 6, we report on the outcomes of our performance analysis. In Section 7, we overview the related work. We conclude in Section 8.

2 Background

Data types and data structures. A sequential *data type* is defined via a set of operations, a set of responses, a set of states, an initial state and a set of transitions. Each transition maps a state and an operation to a new state and a response. A *sequential implementation* (or *sequential data structure*) corresponding to a given data type specifies, for each operation, a sequential read-write algorithm, so that the specification of the data type is respected in every sequential execution.

We consider a system of n asynchronous *processes* (processors or threads of computation) that communicate by performing primitive operations on shared *base objects*. The primitive operations can be reads, writes, or *conditional* operations, such as test&set or compare&swap. A *concurrent implementation* (or *concurrent data structure*) of a given data type assigns, for each process and each operation of the data type, a state machine that is triggered whenever the process invokes an operation and specifies the sequence of *steps* (primitives on the base objects) the process needs to perform to complete the opera-

tion. We require the implementations to be *linearizable* with respect to the data type, i.e., we require that operations *take effect* instantaneously within their intervals [24].

Batched data structures. A *batched implementation* (or *batched data structure*) of a data type exports one function `apply`. This operation takes a *batch* (set) of data type operations as a parameter and returns responses for these operations that are consistent with some sequential application of the operations to the current state of the data structure, which is updated accordingly. We also consider extensions of the definition where we explicitly define the “batched” data type via the *set* [30] or *interval* [10] linearizations. Such a data type takes a batch and a state, and returns a new state and a vector of responses.

For example, in the simplest form, a batched implementation may sequentially apply operations from a batch to the sequential data structure. But batched implementations may also use parallelism to accelerate the execution of the batch: we call these *parallel* batched implementations. We consider two types of parallel batched implementations: *static-multithreading* ones and *dynamic-multithreading* ones [11].

Static multithreading. A parallel batched data structure specifies a distinct (sequential) code to each process in PRAM-like models (PRAM [27], Bulk synchronous parallel model [37], Asynchronous PRAM [18], etc.). For example, in this paper, we provide a batched implementation of a priority queue in the Asynchronous PRAM model. The Asynchronous PRAM consists of n sequential processes, each with its own private local memory, communicating through the shared memory. Each process has its own program. Unlike the classical PRAM model, each process executes its instructions independently of the timing of the other processors. Each process performs one of the four types of instructions per tick of its local clock: global read, global write, local operation, or synchronization step. A synchronization step for a set S of processes is a logical point where each processor in S waits for all the processes in S to arrive before continuing its local program.

Dynamic multithreading. Here the parallel batched implementation is written as a sequential read-write algorithm using concurrency keywords specifying logical parallelism, such as *fork*, *join* and *parallel-for* [11]. An execution of a batch can be presented as a directed acyclic graph (*DAG*) that unfolds dynamically. In the DAG, nodes represent unit-time sequential subcomputations, and edges represent control-flow dependencies between nodes. A node that corresponds to a “fork” has two or more outgoing edges and a node that corresponds to a “join” has two or more incoming edges. The batch is executed using a *scheduler* that chooses which DAG nodes to execute on each process. It can only execute *ready* nodes: not yet executed nodes whose predecessors have all been executed. The most commonly used *work-stealing* scheduler (e.g., [5]) operates as follows. Each process p is provided with a deque for ready nodes. When process p completes node u , it traverses successors of u and collects the ready ones. Then p selects one of the ready successors for execution and adds the remaining ready successors at the bottom of its deque. When p ’s deque is empty, it becomes a *thief*: it randomly picks a victim processor and steals from the top of the victim’s deque.

3 Parallel Combining

In this section, we describe the *parallel combining* technique in a parameterized form: the parameters are specified depending on the application. We then discuss how to use the technique in transforming parallel batched programs into concurrent data structures.

3.1 Combining Data Structure

Our technique relies on a *combining* data structure \mathbb{C} (e.g., the one used in [21]) that maintains a set of requests to a data structure and determines which process is a combiner. If the set of requests is not empty then exactly one process should be a combiner.

Elements stored in \mathbb{C} are of `Request` type consisting of the following fields: 1) the method to be called and its input; 2) the response field; 3) the status of the request with a value in an application-specific `STATUS_SET`; 4) application-specific auxiliary fields. In our applications, `STATUS_SET` contains, by default, values `INITIAL` and `FINISHED`: `INITIAL` meaning that the request is in the initial state, and `FINISHED` meaning that the request is served.

\mathbb{C} supports three operations: 1) `addRequest(r : Request)` inserts request `r` into the set, and the response indicates whether the calling process becomes a combiner or a client; 2) `getRequests()` returns a non-empty set of requests; and 3) `release()` is issued by the combiner to make \mathbb{C} find another process to be a combiner.

In the following, we use any black-box implementation of \mathbb{C} providing this functionality [31, 21, 15, 16].

3.2 Specifying Parameters

To perform an operation, a process executes the following steps (Figure 3): 1) it prepares a request, inserts the request into \mathbb{C} using `addRequest(.)`; 2) if the process becomes the combiner (i.e., `addRequest(.)` returned `true`), it collects requests from \mathbb{C} using `getRequests()`, then it executes algorithm `COMBINER_CODE`, and, finally, it calls `release()` to enable another active process to become a combiner; 3) if the process is a client (i.e., `addRequest(.)` returned `false`), it waits until the status of the request becomes not `INITIAL` and, then, executes algorithm `CLIENT_CODE`.

To use our technique, one should therefore specify `COMBINER_CODE`, `CLIENT_CODE`, and appropriately modify `Request` type and `STATUS_SET`.

Note that *sequential combining* [31, 21, 16, 14] is a special case of parallel combining in which all the work is done by the combiner, and the client code is empty.

3.3 Parallel Batched Algorithms

We discuss how to build a concurrent data structure given a parallel batched one in one of two forms: for static or dynamic multithreading.

In the static multithreading case, each process is provided with a distinct version of `apply` function. We enrich `STATUS_SET` with `STARTED`. In `COMBINER_CODE`, the combiner collects the requests, sets their status to `STARTED`, performs the code of `apply` and waits for the clients to become `FINISHED`. In `CLIENT_CODE` the client waits until its request has `STARTED` status, performs the code of `apply` and sets the status of its request to `FINISHED`.

Suppose that we are given a parallel batched implementation for dynamic multithreading. One can turn it into a concurrent one using parallel combining with the *work-stealing* scheduler. Again, we enrich `STATUS_SET` with `STARTED`. In `COMBINER_CODE`, the combiner collects the requests and sets their status to `STARTED`. Then the combiner creates a working deque, puts there a new node of computational DAG with `apply` function and starts the work-stealing routine on processes-clients. Finally, the combiner waits for the clients to become `FINISHED`. In `CLIENT_CODE`, the client creates a working deque and starts the work-stealing routine.

In Section 5, we illustrate the use of parallel combining and parallel batched programs on the example of a priority queue.

```

1 Request:
2   method
3   input
4   res
5   status ∈ STATUS_SET
6   ...
7
8 execute(method, input):
9   req ← new Request()
10  req.method ← method
11  req.input ← input
12  req.status ← INITIAL
13  if C.addRequest(req):
14    // combiner
15    A ← C.getRequests()
16    COMBINER_CODE
17    C.release()
18  else:
19    while req.status = INITIAL:
20      nop
21    CLIENT_CODE
22  return

```

■ **Figure 3** Parallel combining: pseudocode

4 Read-Optimized Concurrent Data Structures

Before discussing parallel batched algorithms, let us consider a natural application of parallel combining: data structures optimized for *read-dominated* workloads.

Suppose that we are given a sequential data structure D that supports *read-only* (not modifying the data structure) operations, the remaining operations are called *updates*. We assume a scenario where read-only operations dominate over other updates.

Now we explain how to set parameters of parallel combining for this application. At first, `STATUS_SET` consists of three elements `INITIAL`, `STARTED` and `FINISHED`. `Request` type does not have auxiliary fields.

In `COMBINER_CODE` (Figure 4 Lines 1-19), the combiner iterates through the set of collected requests A : if a request contains an update operation then the combiner executes it and sets its status to `FINISHED`; otherwise, the combiner adds the request to set R . Then the combiner sets the status of requests in R to `STARTED`. After that the combiner checks whether its own request is read-only. If so, it executes the method and sets the status of its request to `FINISHED`. Finally, the combiner waits until the status of the requests in R become `FINISHED`.

In `CLIENT_CODE` (Figure 4 Lines 21-24), the client checks whether its method is read-only. If so, the client executes the method and sets the status of the request to `FINISHED`.

► **Theorem 1.** *Algorithm in Figure 4 produces a linearizable concurrent data structure from a sequential one.*

Proof. Any execution of the algorithm can be viewed as a series of non-overlapping combining phases (Figure 4, Lines 2-19). We can group the operations into batches by the combining phase in which they are applied.

```

1 COMBINER_CODE:
2   R ← ∅
3
4   for r ∈ A:
5     if isUpdate(r.method):
6       apply(D, r.method, r.input)
7       r.status ← FINISHED
8     else:
9       R ← R ∪ r
10
11  for r ∈ R:
12    r.status ← STARTED
13  if req.status = STARTED:
14    apply(D, req.method, req.input)
15    req.status ← FINISHED
16
17  for r ∈ R:
18    while r.status = STARTED:
19      nop
20
21 CLIENT_CODE:
22 if not isUpdate(req.method):
23   apply(D, req.method, req.input)
24   req.status ← FINISHED

```

■ **Figure 4** Parallel combining in application to read-optimized data structures

Each update operation is linearized at the point when the combiner applies this operation. Note that this is a correct linearization since all operations that are linearized before are already applied: the operations from preceding combining phases were applied during the preceding phases, while the operations from the current combining phase are applied sequentially by the combiner.

Each read-only operation is linearized at the point when the combiner sets the status of the corresponding request to `STARTED`. By the algorithm, a read-only operation observes all update operations that are applied before and during the current combining phase. Thus, the chosen linearization is correct. ◀

To evaluate the approach in practice we implement a concurrent *dynamic graph* data structure by Holm et al. and execute it in read-dominated environments [25] (Section 6.1).

5 Priority Queue

Priority queue is an abstract data type that maintains an ordered multiset and supports two operations:

- `Insert(v)` — inserts value v into the set;
- $v \leftarrow \text{ExtractMin}()$ — extracts the smallest value from the set.

To the best of our knowledge, no prior parallel implementation of a priority-queue [32, 13, 9, 33] can be efficiently used in our context: their complexity inherently depends on the total number of processes in the system, regardless of the actual batch size. We therefore

introduce a novel heap-based parallel batched priority-queue implementation in a form of `COMBINER_CODE` and `CLIENT_CODE` convenient for parallel combining. The concurrent priority queue is then derived from the described below parallel batched one using the approach presented in Section 3.3.

Here we give only the brief overview of our parallel batched algorithm. Please refer to Appendix A for a detailed description.

In Section 6.2, we show that the resulting concurrent priority queue is able to outperform manually crafted state-of-the-art implementations.

5.1 Sequential Binary Heap

Our batched priority queue is based on the *sequential binary heap* by Gonnet and Munro [19], one of the simplest and fastest sequential priority queues. We briefly describe this algorithm below.

A binary heap of size m is represented as a complete binary tree with nodes indexed by $1, \dots, m$. Each node v has at most two children: $2v$ and $2v + 1$ (to exist, $2v$ and $2v + 1$ should be less than or equal to m). For each node, the *heap property* should be satisfied: the value stored at the node is less than the values stored at its children.

The heap is represented with *size* m and an array a where $a[v]$ is the value at node v . Operations `ExtractMin` and `Insert` are performed as follows:

- `ExtractMin` records the value $a[1]$ as a response, copies $a[m]$ to $a[1]$, decrements m and performs the *sift down* procedure to restore the heap property. Starting from the root, for each node v on the path, we check whether value $a[v]$ is less than values $a[2v]$ and $a[2v + 1]$. If so, then the heap property is satisfied and we stop the operation. Otherwise, we choose the child c , either $2v$ or $2v + 1$, with the smallest value, swap values $a[v]$ and $a[c]$, and continue with c .
- `Insert(x)` initializes a variable val to x , increments m and traverses the path from the root to a new node m . For each node v on the path, if $val < a[v]$, then the two values are swapped. Then the operation continues with the child of v that lies on the path from v to node m . Reaching node m the operation sets its value to val .

The complexity is $O(\log m)$ steps per operation.

5.2 Setup

The heap is defined by its size m and an array a of Node objects. Node object has two fields: value val and boolean $locked$ (In Appendix A it has an additional field).

`STATUS_SET` consists of three items: `INITIAL`, `SIFT` and `FINISHED`.

A Request object consists of: a method $method$ to be called and its input argument v ; a result res field; a $status$ field and a node identifier $start$.

5.3 ExtractMin Phase

Combiner: ExtractMin preparation The combiner withdraws requests A from combining data structure \mathbb{C} . It splits A into sets E and I : the set of `ExtractMin` requests and `Insert` requests. Then it finds $|E|$ nodes $v_1, \dots, v_{|E|}$ of heap with the smallest values using the Dijkstra-like algorithm in $O(|E| \cdot \log |E|)$ steps: (i) create a heap of nodes ordered by values, put there the root 1; (ii) at each of the next $|E|$ steps withdraw the node v with the minimal value from the heap; (iii) put two children of v , $2v$ and $2v + 1$, to the heap. The $|E|$ withdrawn nodes are the nodes with the $|E|$ minimal values. For each request $E[i]$, the combiner sets $E[i].res$ to $a[v_i].val$, $a[v_i].locked$ to `true`, and $E[i].start$ to v_i .

The combiner proceeds by pairing Insert requests in I with ExtractMin requests in E using the following procedure. Suppose that $\ell = \min(|E|, |I|)$. For each $i \in [1, \ell]$, the combiner sets $a[v_i].val$ to $I[i].v$ and $I[i].status$ to **FINISHED**, i.e., this Insert request becomes completed. Then, for each $i \in [\ell + 1, |E|]$, the combiner sets $a[v_i].val$ to the value of the last node $a[m]$ and decrements m , as in the sequential algorithm. Finally, the combiner sets the status of all requests in E to **SIFT**.

Clients: ExtractMin phase Briefly, the clients sift down the values in nodes $v_1, \dots, v_{|E|}$ in parallel using hand-over-hand locking: the *locked* field of a node is set whenever there is a *sift down* operation working on that node.

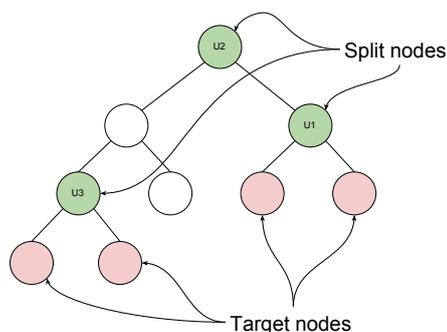
A client c waits until the status of its request becomes **SIFT**. c starts sifting down from $req.start$. Suppose that c is currently at node v . c waits until the *locked* fields of the children become **false**. If $a[v].val$, the value of v , is less than the values in its children, then *sift down* is finished: c unsets $a[v].locked$ and sets the status of its request to **FINISHED**. Otherwise, let w be the child with the smallest value. Then c swaps $a[v].val$ and $a[w].val$, sets $a[w].locked$, unsets $a[v].locked$ and continues with node w .

If the request of the combiner is ExtractMin, it also runs the code above as a client. The combiner considers the ExtractMin phase completed when all requests in E have status **FINISHED**.

5.4 Insert Phase

For simplicity, we describe first the sequential algorithm.

At first, the combiner removes all completed requests from I . Then it initializes new nodes $m + 1, \dots, m + |I|$ which we call *target nodes* and increments m by $|I|$. The nodes for which the subtrees of both children contain at least one target node are called *split nodes*. (See Figure 5 for an example of how target and split nodes can be defined.)



■ **Figure 5** Split and target nodes

The combiner collects the values of the remaining Insert requests and sorts them: $r_1, \dots, r_{|I|}$. Then it sets the status of these requests to **FINISHED**.

Now, we introduce InsertSet class: it consists of two sorted lists A and B . The combiner starts the following recursive procedure at the root with InsertSet s : $s.A$ contains $r_1, \dots, r_{|I|}$ while $s.B$ is empty. Suppose that the procedure is called on node v and InsertSet s . Let min be the minimum out of the first element of $s.A$ and the first element of $s.B$. If v is a target node then the combiner sets $a[v].res$ to min and withdraws m from the corresponding list. Otherwise, the combiner compares $a[v].res$ with

min : if $a[v].res$ is smaller, then it does nothing; otherwise, it appends $a[v].res$ to the end of $s.B$ (note that $s.B$ remains sorted because $s.B$ consists only of values that were ancestors in the heap), withdraws min from the corresponding list and sets $a[v].res$ to min .

If v is not the split node the combiner calls the recursive procedure on the child with target nodes in the subtree and with InsertSet s . Otherwise, the combiner calculates inL and inR — the number of target nodes in the left and right subtrees of v . Suppose, for

simplicity, that inL is less than inR (the opposite case can be resolved similarly). The combiner splits s into two parts: create `InsertSet` s_L , move $\min(inL, |s.A|)$ first values from $s.A$ to $s_L.A$ and move $\min(inL - |s_L.A|, |s.B|)$ first values from $s.B$ to $s_L.B$. Finally, it calls the recursive procedure on the left child with `InsertSet` s_L and on the right child with `InsertSet` s .

This algorithm works in $O(\log m + c \log c)$ steps (and can be optimized to $O(\log m + c)$ steps), where m is the size of the queue and c is the number of `Insert` requests to apply. Note that this algorithm is almost non-parallelizable due to its small complexity, and our parallel algorithm is only developed to reduce constant factors.

Now, we construct a parallel algorithm for the `Insert` phase. We enrich `Node` object with the `IntegerSet` field *split*. The combiner sets the *start* field of the first client ($i[1].start$) to the root 1, while *start* fields of other clients to the right children of split nodes (we have exactly $|I| - 1$ split nodes). Then it initializes the *split* field of the root as the `IntegerSet` s is initialized at the beginning of the sequential algorithm: list A contains values of requests while list B is empty.

Each client waits until the *split* field of the corresponding *start* node is non-null. Then it reads this `IntegerSet`: the values from this set should be inserted in the subtree. Finally, the client performs the procedure similar to the recursive procedure from the sequential algorithm except for one difference: when it reaches a split node instead of going recursively to left and right children, it splits `InsertSet` to s_L and s_R of sizes inL and inR , puts s_R into the *split* field of the right child (in order to wake another client) and continues with the left child and s_L .

For further details about the parallel algorithm we refer to Appendix A.

6 Experiments

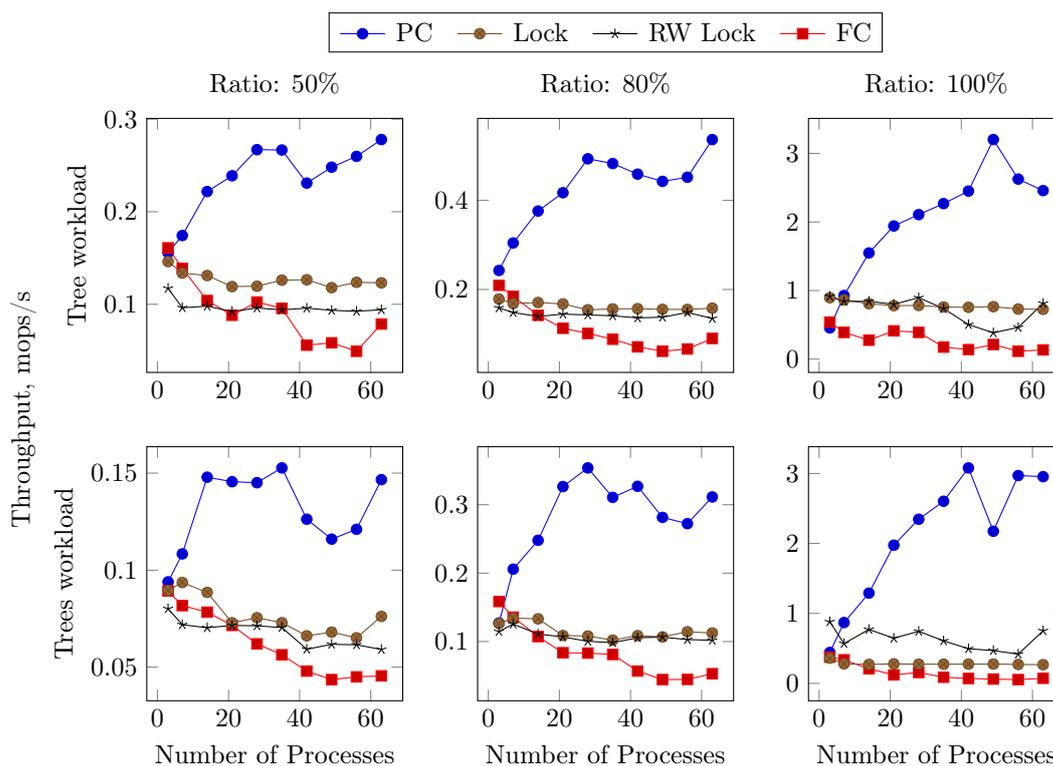
We evaluate Java implementations of our data structures on a 4-processor AMD Opteron 6378 2.4 GHz server with 16 threads per processor (yielding 64 threads in total), 512 Gb of RAM, running Ubuntu 14.04.5 with Java 1.8.0_111-b14 and HotSpot JVM 25.111-b14.

6.1 Concurrent Dynamic Graph

To illustrate how parallel combining can be used to construct read-optimized concurrent data structures, we took the sequential dynamic graph implementation by Holm et al. [25]. This data structure supports two update methods: an insertion of an edge and a deletion of an edge; and one read-only method: a connectivity query that tests whether two vertices are connected.

We compare our implementation based on parallel combining (PC) with *flat combining* [21] as a combining data structure against three others: (1) Lock, based on `ReentrantLock` from `java.util.concurrent`; (2) RW Lock, based on `ReentrantReadWriteLock` from `java.util.concurrent`; and (3) FC, based on *flat combining* [21]. The code is available at <https://github.com/Aksenov239/concurrent-graph>.

We consider workloads parametrized with: 1) the fraction x of connectivity queries (50%, 80% or 100%, as we consider read-dominated workloads); 2) the set of edges E : edges of a single random tree, or edges of ten random trees; 3) the number of processes P (from 1 to 64). We prepopulate the graph on 10^5 vertices with edges from E : we insert each edge with probability $\frac{1}{2}$. Then we start P processes. Each process repeatedly performs operations: 1) with probability x , it calls a connectivity query on two vertices chosen uniformly at



■ **Figure 6** Dynamic graph implementations

random; 2) with probability $1 - \frac{\alpha}{2}$, it inserts an edge chosen uniformly at random from E ; 3) with probability $1 - \frac{\alpha}{2}$, it deletes an edge chosen uniformly at random from E .

We denote the workloads with E as a single tree as *Tree* workloads, and other workloads as *Trees* workloads. *Tree* workloads are interesting because they show the degenerate case: the dynamic graph behaves as a dynamic tree. In this case, about 50% of update operations successfully change the spanning forest, while other update operations only check the existence of the edge and do not modify the graph. *Trees* workloads are interesting because a reasonably small number (approximately, 5-10%) of update operations modify the maintained set of all edges and the underlying complex data structure that maintains a spanning forest (giving in total the squared logarithmic complexity), while other update operations can only modify the set of edges but cannot modify the underlying complex data structure (giving in total the logarithmic complexity).

For each setting and each algorithm, we run the corresponding workload for 10 seconds to warmup HotSpot JVM and then we run the workload five more times for 10 seconds. The average throughput of the last five runs is reported in Figure 6.

From the plots we can infer two general observations: PC exhibits the highest throughput over all considered implementations and it is the only one whose throughput scales up with the number of the processes. On the 100% workload we expect the throughput curve to be almost linear since all operations are read-only and can run in parallel. The plots almost confirm our expectation: the curve of the throughput is a linear function with coefficient $\frac{1}{2}$ (instead of the ideal coefficient 1). We note that this is almost the best we can achieve: a combiner typically collects operations of only approximately half the number of working processes. In addition, the induced overhead is still perceptible, since each connectivity query

works in just logarithmic time. With the decrease of the fraction of read-only operations we expect that the throughput curve becomes flatter, as plots for the 50% and 80% workloads confirm.

It is also interesting to point out several features of other implementations. At first, FC implementation works slightly worse than Lock and RW Lock. This might be explained as follows. Lock implementations (`ReentrantLock` and `ReentrantReadWriteLock`) behind Lock and RWLock implementations are based on CLH Lock [12] organized as a *queue*: every competing process is appended to the queue and then waits until the previous one releases the lock. Operations on the dynamic graph take significant amount of time, so under high load when the process finishes its operation it appends itself to the queue in the lock without any contention. Indeed, all other processes are likely to be in the queue and, thus, no process can contend. By that the operations by processes are serialized with almost no overhead. In contrast, the combining procedure in FC introduces non-negligible overhead related to gathering the requests and writing them into requests structures.

Second, it is interesting to observe that, against the intuition, RWLock is not so superior with respect to Lock on read-only workloads. As can be seen, when there are update operations in the workload RWLock works even worse than Lock. We relate this to the fact that the overhead hidden inside `ReentrantReadWriteLock` spent on manipulation with read and write requests is bigger than the overhead spent by `ReentrantLock`. With the increase of the percentage of read-only operations the difference between Lock and RWLock diminishes and RWLock becomes dominant since read operations become more likely to be applied concurrently (for example, on 50% it is normal to have an execution without any parallelization: read operation, write operation, read operation, and so on). However, on 100% one could expect that RWLock should exhibit ideal throughput. Unfortunately, in this case, under the hood `ReentrantReadWriteLock` uses compare&swap on the shared variable that represents the number of current read operations. Read-only operations take enough time but not enough to amortize the considerable traffic introduced by concurrent compare&swaps. Thus, the plot for RWLock is almost flat, getting even slightly worse with the increase of the number of processes, and we blame the traffic for this.

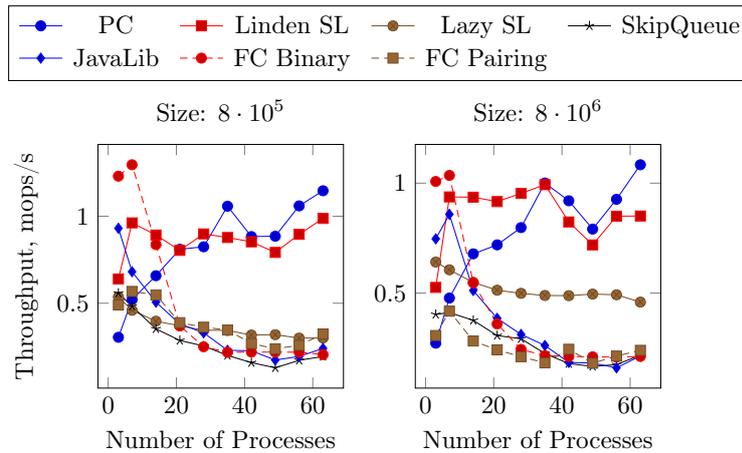
6.2 Priority Queue

We run our algorithm (PC) with *flat combining* [21] as a combining data structure against six state-of-the-art concurrent priority queues: (1) the lock-free skip-list by Linden and Johnson (Linden SL [28]), (2) the lazy lock-based skip-list (Lazy SL [23]), (3) the non-linearizable lock-free skip-list by Herlihy and Shavit (SkipQueue [23]) as an adaptation of Lotan and Shavit's algorithm [34], (4) the lock-free skip-list from Java library (JavaLib), (5) the binary heap with flat combining (FC Binary [21]), and (6) the pairing heap with flat combining (FC Pairing [21]).¹ The code is available at <https://github.com/Aksenov239/FC-heap>.

We consider workloads parametrized by: 1) the initial size of the queue S ($8 \cdot 10^5$ or $8 \cdot 10^6$); and 2) the number P of working processes (from 1 to 64). We prepopulate the queue with S random integers chosen uniformly from the range $[0, 2^{31} - 1]$. Then we start P processes, and each process repeatedly performs operations: with equal probability it either inserts a random value taken uniformly from $[0, 2^{31} - 1]$ or extracts the minimum value.

For each setting and each algorithm, we run the corresponding workload for 10 seconds

¹ We are aware of the cache-friendly priority queue by Braginsky et al. [8], but we do not have its Java implementation.



■ **Figure 7** Priority Queue implementations

to warmup HotSpot JVM and then we run the workload five more times for 10 seconds. The average throughput of the last five runs is reported in Figure 7.

On a small number of processes (< 15), PC performs worse than other algorithms. With respect to Linden SL, Lazy SL, SkipQueue and JavaLib this can be explained by two different issues:

- Synchronization incurred by PC is not compensated by the work done;
- Typically, a combiner collects operations of only approximately half the processes, thus, we “overserialize”, i.e., only $\frac{n}{2}$ operations can be performed in parallel.

In contrast, on small number of processes, the other four algorithms can perform operations almost with no contention. With respect to algorithms based on flat combining, FC Binary and FC Pairing, our algorithm is simply slower on one process than the simplest sequential binary and pairing heap algorithms.

With the increase of the number of processes the synchronization overhead significantly increases for all algorithms (in addition to the fact that FC Binary and FC Pairing cannot scale). As a result, starting from 15 processes, PC outperforms all algorithms except for Linden SL. Linden SL relaxes the contention during ExtractMin operations, and it helps to keep the throughput approximately constant. At approximately 40 processes the benefits of the parallel batched algorithm in PC starts prevailing the costs of explicit synchronization, and our algorithms overtakes Linden SL.

It is interesting to note that FC Binary performs very well when the number of processes is small: the overhead on the synchronization is very small, the processes are from the same core and the simplest binary heap performs operations very fast.

7 Related Work

To the best of our knowledge, Yew et al. [38] were the first to propose combining concurrent operations. They introduced a *combining tree*: processes start at distinct leaves, traverse upwards, and gain exclusive access by reaching the root. If, during the traversal, two processes access the same tree node, one of them adopts the operations of another and continues the traversal, while the other waits until its operations are completed. Several improvements of this technique have been discussed, such as adaptive combining tree [36], barrier implementations [20, 29] and counting networks [35].

A different approach was proposed by Oyama et al. [31]. Here the data structure is protected by a lock. A thread with a new operation to be performed adds it to a list of submitted requests and then tries to acquire the lock. The winner of the lock performs the pending requests on behalf of other processes from the list in LIFO order. The main drawback of this approach is that all processes have to perform CAS on the head of the list. The *flat combining* technique presented by Hendler et al. [21] addresses this issue by replacing the list of requests with a *publication list* which maintains a distinct *publication record* per participating process. A process puts its new operation in its publication record, and the publication record is only maintained in the list if the process is “active enough”. This way the processes generally do not contend on the head of the list. Variations of flat combining were later proposed for various contexts [15, 16, 26, 14].

Hierarchical combining [22] is the first attempt to improve performance of combining using the computational power of clients. The list of requests is split into blocks, and each of these blocks has its own combiner. The combiners push the combined requests from the block into the second layer implemented as the standard flat combining with one combiner. This approach, however, may be sub-optimal as it does not involve *all* clients. Moreover, this approach works only for specific data structures, such as stacks or unfair synchronous queues, where operations could be combined without accessing the data structure.

In a different context, Agrawal et al. [2] suggested to use a *parallel batched* data structure instead of a concurrent one. They provide provable bounds on the running time of a dynamic multithreaded parallel program using P processes and a specified *scheduler*. The proposed scheduler extends the work-stealing scheduler by maintaining separate *batch* work-stealing deques that are accessed whenever processes have operations to be performed on the abstract data type. A process with a task to be performed on the data structure stores it in a *request array* and tries to acquire a global lock. If succeeded, the process puts the task to perform the batch update in its batch deque. Then all the processes with requests in the request array run the work-stealing routine on the batch deques until there are no tasks left. The idea of [2] is similar to ours. However, their algorithm is designed for systems with the fixed set of processes, whereas we allow the processes to join and leave the execution. From a more formal perspective, our goals are different: we aim at improving the performance of a *concurrent data structure* while their goal was to establish bounds on the running time of a *parallel program in dynamic multithreading*. Furthermore, implementing a concurrent data structure from its parallel batched counterpart for dynamic multithreading is only one of the applications of our technique, as sketched in Section 3.3.

8 Concluding remarks

Besides performance gains, parallel combining can potentially bring other interesting benefits.

First, a parallel batched implementation is typically provided with bounds on the running time. The use of parallel combining might allow us to derive bounds on the operations of resulting *concurrent* data structures. Consider, for example, a binary search tree. To balance the tree, state-of-the-art concurrent algorithms use the relaxed AVL-scheme [7]. This scheme guarantees that the height of the tree never exceeds the contention level (the number of concurrent operations) plus the logarithm of the tree size. Applying parallel combining to a parallel batched binary search tree (e.g., [4]), we get a concurrent tree with a strict logarithmic bound on the height.

Second, the technique might enable the first ever concurrent implementation of certain

data types, for example, a *dynamic tree* [1].

As shown in Section 6, our concurrent priority queue performs well compared to state-of-the-art algorithms. A possible explanation is that the underlying parallel batched implementation is designed for *static* multithreading and, thus, it has little synchronization overhead. This might not be the case for implementations based on *dynamic* multithreading, where the overhead induced by the scheduler can be much higher. We intend to explore this distinction in the forthcoming work.

References

- 1 Umut A. Acar, Vitaly Aksenov, and Sam Westrick. Brief announcement: Parallel dynamic tree contraction via self-adjusting computation. In *SPAA*, pages 275–277. ACM, 2017.
- 2 Kunal Agrawal, Jeremy T Fineman, Kefu Lu, Brendan Sheridan, Jim Sukha, and Robert Utterback. Provably good scheduling for parallel programs that use data structures through implicit batching. In *SPAA*, pages 84–95. ACM, 2014.
- 3 Timo Bingmann, Thomas Keh, and Peter Sanders. A bulk-parallel priority queue in external memory with stxxl. In *International Symposium on Experimental Algorithms*, pages 28–40. Springer, 2015.
- 4 Guy E Blelloch, Daniel Ferizovic, and Yihan Sun. Just join for parallel ordered sets. In *SPAA*, pages 253–264. ACM, 2016.
- 5 Robert D Blumofe and Charles E Leiserson. Scheduling multithreaded computations by work stealing. *Journal of the ACM (JACM)*, 46(5):720–748, 1999.
- 6 OpenMP Architecture Review Board. Openmp application interface. <http://www.openmp.org>, 2008.
- 7 Luc Bougé, Joaquim Gabarro, Xavier Messeguer, Nicolas Schabanel, et al. Height-relaxed AVL rebalancing: A unified, fine-grained approach to concurrent dictionaries. Technical report, ENS Lyon, 1998.
- 8 Anastasia Braginsky, Nachshon Cohen, and Erez Petrank. Cbpq: High performance lock-free priority queue. In *Euro-Par*, pages 460–474. Springer, 2016.
- 9 Gerth Stølting Brodal, Jesper Larsson Träff, and Christos D Zaroliagis. A parallel priority queue with constant time operations. *Journal of Parallel and Distributed Computing*, 49(1):4–21, 1998.
- 10 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. Specifying concurrent problems: Beyond linearizability and up to tasks. In *DISC*, pages 420–435, 2015.
- 11 Thomas H Cormen, Charles Eric Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. The MIT press, 3rd edition, 2009.
- 12 Travis Craig. Building fifo and priorityqueuing spin locks from atomic swap. Technical report, Technical Report TR 93-02-02, University of Washington, 02 1993., 1993.
- 13 Narsingh Deo and Sushil Prasad. Parallel heap: An optimal parallel priority queue. *The Journal of Supercomputing*, 6(1):87–98, 1992.
- 14 Dana Drachler-Cohen and Erez Petrank. Lcd: Local combining on demand. In *OPODIS*, pages 355–371. Springer, 2014.
- 15 Panagiota Fatourou and Nikolaos D Kallimanis. A highly-efficient wait-free universal construction. In *SPAA*, pages 325–334. ACM, 2011.
- 16 Panagiota Fatourou and Nikolaos D Kallimanis. Revisiting the combining synchronization technique. In *ACM SIGPLAN Notices*, volume 47, pages 257–266. ACM, 2012.
- 17 Matteo Frigo, Charles E Leiserson, and Keith H Randall. The implementation of the cilk-5 multithreaded language. *ACM Sigplan Notices*, 33(5):212–223, 1998.
- 18 Phillip B Gibbons. A more practical pram model. In *SPAA*, pages 158–168. ACM, 1989.

- 19 Gaston H Gonnet and J Ian Munro. Heaps on heaps. *SIAM Journal on Computing*, 15(4):964–971, 1986.
- 20 Rajiv Gupta and Charles R Hill. A scalable implementation of barrier synchronization using an adaptive combining tree. *International Journal of Parallel Programming*, 18(3):161–180, 1989.
- 21 Danny Hendler, Itai Incze, Nir Shavit, and Moran Tzafrir. Flat combining and the synchronization-parallelism tradeoff. In *SPAA*, pages 355–364, 2010.
- 22 Danny Hendler, Itai Incze, Nir Shavit, and Moran Tzafrir. Scalable flat-combining based synchronous queues. In *DISC*, pages 79–93. Springer, 2010.
- 23 Maurice Herlihy and Nir Shavit. *The art of multiprocessor programming*. Morgan Kaufmann, 2012.
- 24 Maurice Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- 25 Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *Journal of the ACM (JACM)*, 48(4):723–760, 2001.
- 26 Brandon Holt, Jacob Nelson, Brandon Myers, Preston Briggs, Luis Ceze, Simon Kahan, and Mark Oskin. Flat combining synchronized global data structures. In *7th International Conference on PGAS Programming Models*, page 76.
- 27 Joseph JáJá. *An introduction to parallel algorithms*, volume 17. Addison-Wesley Reading, 1992.
- 28 Jonatan Lindén and Bengt Jonsson. A skiplist-based concurrent priority queue with minimal memory contention. In *International Conference On Principles Of Distributed Systems*, pages 206–220. Springer, 2013.
- 29 John M Mellor-Crummey and Michael L Scott. Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Transactions on Computer Systems (TOCS)*, 9(1):21–65, 1991.
- 30 Gil Neiger. Set-linearizability. In *PODC*, page 396, 1994.
- 31 Yoshihiro Oyama, Kenjiro Taura, and Akinori Yonezawa. Executing parallel programs with synchronization bottlenecks efficiently. In *Proceedings of the International Workshop on Parallel and Distributed Computing for Symbolic and Irregular Applications*, volume 16, 1999.
- 32 Maria Cristina Pinotti and Geppino Pucci. Parallel priority queues. *Information Processing Letters*, 40(1):33–40, 1991.
- 33 Peter Sanders. Randomized priority queues for fast parallel access. *Journal of Parallel and Distributed Computing*, 49(1):86–97, 1998.
- 34 Nir Shavit and Itay Lotan. Skiplist-based concurrent priority queues. In *IPDPS*, pages 263–268. IEEE, 2000.
- 35 Nir Shavit and Asaph Zemach. Diffracting trees. *ACM Transactions on Computer Systems (TOCS)*, 14(4):385–428, 1996.
- 36 Nir Shavit and Asaph Zemach. Combining funnels: a dynamic approach to software combining. *Journal of Parallel and Distributed Computing*, 60(11):1355–1387, 2000.
- 37 Leslie G Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- 38 Pen-Chung Y, Nian-Feng T, et al. Distributing hot-spot addressing in large-scale multiprocessors. *IEEE Transactions on Computers*, 100(4):388–395, 1987.

A Priority Queue with Parallel Combining

Priority queue is an abstract data type that maintains an ordered multiset and supports two operations:

- `Insert(v)` — inserts value v into the set;
- $v \leftarrow \text{ExtractMin}()$ — extracts the smallest value from the set.

A.1 Batched Priority Queues: Review

Several batched priority queues were proposed in the literature for different parallel machines [27].

Pinotti and Pucci [32] proposed a batched priority queue for a p -processor CREW PRAM implemented as a heap each node of which contains p values in sorted order: only batches of size p are accepted.

Deo and Prasad [13] proposed a similar batched priority queue for a p -processor EREW PRAM which can accept batches of any size not exceeding p . But the batch processing time is proportional to p . For example, even if the batch consists of only one operations and only one process is involved in the computation, the execution still takes $O(p \log m)$ time on a queue of size m .

Brodal et al. [9] proposed a batched priority queue that accepts batches of `Insert` and `DecreaseKey` operations, but not the batches of `ExtractMin` operations. The priority queue maintains a set of pairs (key, element) ordered by keys. `Insert` operation takes a pair (k, e) and inserts it. `DecreaseKey` operations takes a pair (d, e) , searches in the queue for a pair (d', e) such that $d < d'$ and then replaces (d', e) with (d, e) .

Sanders [33] developed a randomized distributed priority queue for MIMD. MIMD computer has p processing elements that communicate via asynchronous message-passing. Again, the batch execution time is proportional to p , regardless of the batch size.

Bingmann et al. [3] described a variation of the priority queue by Sanders [33] for external memory and, thus, it has the same issue.

To summarize, all earlier implementations we are aware of are tailored for a fixed number of processes p . As a result, (1) the running time of the algorithms always depend on p , regardless of the batch size and the number of involved processes; (2) once a data structure is constructed, we are unable to introduce more processes into system and use them efficiently. To respond to this issues, we propose a new parallel batched algorithm that applies a batch of size c to a queue of size m in $O(c \cdot (\log c + \log m))$ RMRs for CC or DSM models in total and in $O(c + \log m)$ RMRs per process. By that, our algorithm can use up to $c \approx \log m$ processes efficiently.

A.2 Sequential Binary Heap

Our batched priority queue is based on the *sequential binary heap* by Gonnet and Munro [19], one of the simplest and fastest sequential priority queues. We briefly describe this algorithm below.

A binary heap of size m is represented as a complete binary tree with nodes indexed by $1, \dots, m$. Each node v has at most two children: $2v$ and $2v + 1$ (to exist, $2v$ and $2v + 1$ should be less than or equal to m). For each node, the *heap property* should be satisfied: the value stored at the node is less than the values stored at its children.

The heap is represented with *size* m and an array a where $a[v]$ is the value at node v . Operations `ExtractMin` and `Insert` are performed as follows:

```

1 class InsertSet:
2     List A, B
3
4     split(int l):
5         Pair<InsertSet, InsertSet>
6         L ← min(l, |A| + |B| - 1)
7         X ← new InsertSet()
8         if |A| ≥ L:
9             for i in 1..L:
10                a ← X.A.removeFirst()
11                X.A.append(a)
12        else:
13            for i in 1..L:
14                b ← X.B.removeFirst()
15                X.B.append(b)
16
17        if L = l:
18            return (X, this)
19
20        else:
21            return (this, X)
22
23 class Node:
24     V val
25     bool locked
26     InsertSet split
27
28 class Request:
29     method: { ExtractMin, Insert }
30     V v
31     V res
32     STATUS_SET status
33     int start
34
35     // for client_insert(req)
36     // specifies the segment of leaves
37     // in a subtree of start node
38     int l, r

```

■ **Figure 8** Parallel Priority Queue. Classes

- **ExtractMin** records the value $a[m]$ as a response, copies $a[m]$ to $a[1]$, decrements m and performs the *sift down* procedure to restore the heap property. Starting from the root, for each node v on the path, we check whether value $a[v]$ is less than values $a[2v]$ and $a[2v+1]$. If so, then the heap property is satisfied and we stop the operation. Otherwise, we choose the child c , either $2v$ or $2v+1$, with the smallest value, swap values $a[v]$ and $a[c]$, and continue with c .
- **Insert**(x) initializes a variable val to x , increments m and traverses the path from the root to a new node m . For each node v on the path, if $val < a[v]$, then the two values are swapped. Then the operation continues with the child of v that lies on the path from v to node m . Reaching node m the operation sets its value to val .

The complexity is $O(\log m)$ steps per operation.

A.3 Combiner and Client. Classes

Now, we describe our novel parallel batched priority queue in the form of `COMBINER_CODE` and `CLIENT_CODE` that fits the parallel combining framework described in Section 3. It is based on the sequential binary heap by Gonnet and Munro [19]. The code of necessary classes is presented in Figure 8, `COMBINER_CODE` is presented in Figure 9, and `CLIENT_CODE` is presented in Figure 10.

We introduce a sequential object `InsertSet` (Figure 8 Lines 1-20) that consists of two sorted linked lists A and B supporting *size* operation $|\cdot|$. The size of `InsertSet` S is $|S| = |S.A| + |S.B|$. `InsertSet` supports operation `split`: $(X, Y) \leftarrow S.\text{split}(\ell)$, which splits `InsertSet` S into two `InsertSet` objects X and Y , where $|X| = \ell$ and $|Y| = |S| - \ell$. This operation is executed sequentially in $O(L = \min(\ell, |S| - \ell))$ steps. The split operation works as follows: 1) new `InsertSet` T is created (Line 7); 2) if $|S.A| \geq L$ then the first L values of $S.A$ are moved to $T.A$ (Lines 9-11); otherwise, the first L values from $S.B$ are moved to $T.B$ (Lines 13-15); note that either $|S.A|$ or $|S.B|$ should be at least L ; 3) if $L = \ell$ then (T, S) is

```

1  A ← C.getRequests()
2
3  if m ≤ |A|:
4    apply A sequentially
5    for r ∈ A:
6      r.status ← FINISHED
7    return
8
9  E ← ∅
10 I ← ∅
11
12 for r ∈ A:
13   if isInsert(r):
14     I ← I ∪ r
15   else:
16     E ← I ∪ r
17
18 bestE ← new int[|E|]
19 heap ← new Heap<Pair<V, int>>()
20
21 heap.insert((a[1], 1))
22
23 for i in 1..|E|:
24   (v, id) ← heap.extract_min()
25   bestE[i] ← id
26   heap.insert((a[2 · v], 2 · v))
27   heap.insert(
28     (a[2 · v + 1], 2 · v + 1))
29
30 for i in 1..|E|:
31   E[i].res ← a[bestE[i]].val
32   a[bestE[i]].locked ← true
33   E[i].start ← bestE[i]
34
35 l ← min(|E|, |I|)
36 for i in 1..l:
37   a[bestE[i]] ← I[i].v
38   I[i].status ← FINISHED
39
40 for i in l + 1..|E|:
41   a[bestE[i]] ← a[m]
42   m--
43
44 for i in 1..|E|:
45   E[i].status ← SIFT
46
47 if req.status = SIFT:
48   client_extract_min(req)
49
50 for i in 1..|E|:
51   while E[i].status = SIFT:
52     nop
53 I ← I[1 + 1..|I|]
54
55 I[1].start ← 1
56 I[1].l ← 2⌊log2(m+|I|)⌋
57 I[1].r ← 2 · I[1].l - 1
58 for i in 2..|I|:
59   t ← m + i
60   power ← 1
61   while t > 1:
62     p ← t / 2
63     if 2 · p = t:
64       break
65     t ← p
66     power ← 2 · power
67
68   if t = 1:
69     t ← 2
70   I[i].start ← t + 1
71   I[i].l ← I[i].start · power
72   I[i].r ← I[i].l + power - 1
73
74 // L and R are global variables
75 // necessary for client_insert(req)
76 L ← m + 1
77 R ← m + |I|
78
79 m ← m + |I|
80
81 args ← new V[|I|]
82 for i in 1..|I|:
83   args[i] ← I[i].v
84
85 sort(args)
86
87 a[1].split ← new InsertSet()
88 for i in 1..|I|:
89   a[1].split.A.append(args[i])
90
91 for i in 1..|I|:
92   I[i].status ← SIFT
93
94 if req.status = SIFT:
95   client_insert(req)
96
97 for i in 1..|I|:
98   while I[i].status = SIFT:
99     nop

```

■ **Figure 9** Parallel Priority Queue. COMBINER_CODE

```

1 CLIENT_CODE:
2   if isInsert(req):
3     if req.status = SIFT:
4       client_insert(req)
5     else:
6       client_extract_min(req)
7   req.status ← FINISHED
8   return
9
10 client_extract_min(Request req):
11   v ← req.start
12   while 2 · v ≤ m:
13     while a[2 · v].locked:
14       nop
15     if 2 · v + 1 ≤ m:
16       while a[2 · v + 1].locked:
17         nop
18     c ← 2 · v
19     if 2 · v + 1 ≤ m
20       and a[2 · v] > a[2 · v + 1]:
21       c ← 2 · v + 1
22     if a[c] > a[v]:
23       a[v].locked ← false
24       break
25     else:
26       swap(a[c], a[v])
27       a[c].locked ← true
28       a[v].locked ← false
29       v ← c
30   return
31
32 // Integers that specifies
33 // a segment of target nodes,
34 // i.e., m + 1 and m + |I|
35 global int L, R
36
37 targets_in_subtree(l, r): int
38   return min(r, R) - max(l, L) + 1
39
40 client_insert(Request req):
41   v ← req.start
42   while a[v].split = null:
43     nop
44   S ← a[v].split
45   a[v].split ← null
46
47   l ← req.l
48   r ← req.r
49
50   while v ∉ [L, R]:
51     a ← S.A.first()
52     b ← S.B.first()
53     if a[v] < min(a, b):
54       x ← a[v]
55       a[v] ← min(a, b)
56       if a < b:
57         S.A.pollFirst()
58       else:
59         S.B.pollFirst()
60       S.B.append(x)
61
62   mid ← (l + r) / 2
63   inL ←
64     targets_in_subtree(l, mid)
65   inR ←
66     targets_in_subtree(mid + 1, r)
67
68   if inL = 0:
69     v ← 2 · v + 1
70     l ← mid + 1
71
72   if inR = 0:
73     v ← 2 · v
74     r ← mid
75
76   if inL ≠ 0 and inR ≠ 0:
77     (S, T) ← S.split(inL)
78     a[2 · v + 1].split ← T
79     v ← 2 · v
80     r ← mid
81
82   if |S.A| ≠ 0:
83     a[v] ← S.A.first()
84   else:
85     a[v] ← S.B.first()
86   return

```

■ Figure 10 Parallel Priority Queue. CLIENT_CODE

returned; otherwise, (S, T) is returned (Lines 17-20).

The heap is defined by its size m and an array a of Node objects. Node object (Figure 8 Lines 21-24) has three fields: value val , boolean $locked$ and InsertSet $split$.

STATUS_SET consists of three items: INITIAL, SIFT and FINISHED.

A Request object (Figure 8 Lines 26-36) consists of: a method $method$ to be called and its input argument v ; a result res field; a $status$ field; a node identifier $start$.

A.4 ExtractMin Phase

Combiner: ExtractMin preparation (Figure 9 Lines 1-52). First, the combiner withdraws requests A from the combining data structure \mathbb{C} (Line 1). If the size of A is larger than m , the combiner serves the requests sequentially (Lines 3-7). Intuitively, in this case, there is no way to parallelize the execution. For example, if A consists of only Insert requests and if there are more Insert requests than the number of nodes in the corresponding binary tree, we cannot insert them in parallel.

In the following, we assume that the size of the queue is at least the size of A . The combiner splits A into sets E and I (Lines 9-16): the set of ExtractMin requests and the set of Insert requests. Then it finds $|E|$ nodes $v_1, \dots, v_{|E|}$ of heap with the smallest values (Lines 18-28), e.g., using the Dijkstra-like algorithm in $O(|E| \cdot \log |E|)$ primitive steps or $O(|E|)$ RMRs. For each request $E[i]$, the combiner sets $E[i].res$ to $a[v_i].val$, $a[v_i].locked$ to true, and $E[i].start$ to v_i (Lines 30-33).

The combiner proceeds by pairing Insert requests in I with ExtractMin requests in E using the following procedure. (Lines 35-38). Suppose that $\ell = \min(|E|, |I|)$. For each $i \in [1, \ell]$, the combiner sets $a[v_i].val$ to $I[i].v$ and $I[i].status$ to FINISHED, i.e., this Insert request becomes completed. Then, for each $i \in [\ell + 1, |E|]$, the combiner sets $a[v_i].val$ to the value of the last node $a[m]$ and decreases m , as in the sequential algorithm (Lines 40-42). Finally, the combiner sets the status of all requests in E to SIFT (Lines 44-45).

Clients: ExtractMin phase (Figure 10 Lines 11-29). Briefly, the clients sift down the values in nodes $v_1, \dots, v_{|E|}$ in parallel using hand-over-hand locking: the $locked$ field of a node is set whenever there is a *sift down* operation working on that node.

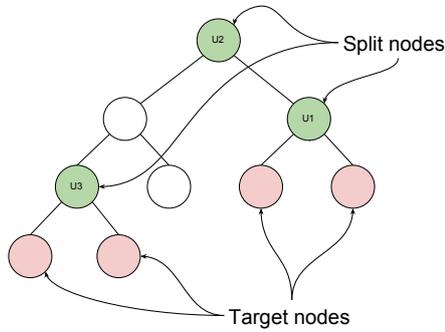
A client c waits until the status of its request becomes SIFT. c starts sifting down from $req.start$. Suppose that c is currently at node v . c waits until the $locked$ fields of the children become false (Lines 13-17). If $a[v].val$, the value of v , is less than the values in its children, then *sift down* is finished (Lines 23-24): c unsets $a[v].locked$ and sets the status of its request to FINISHED. Otherwise, let w be the child with the smallest value. Then c swaps $a[v].val$ and $a[w].val$, sets $a[w].locked$, unsets $a[v].locked$ and continues with node w (Lines 26-29).

If the request of the combiner is ExtractMin, it also runs the code above as a client (Figure 9 Lines 47-48). The combiner considers the ExtractMin phase completed when all requests in E have status FINISHED (Lines 50-52).

A.5 Insert Phase

Combiner: Insert preparation (Figure 9 Lines 53-99). For Insert requests, the combiner removes all completed requests from I (Line 53). Nodes $m + 1, \dots, m + |I|$ have to be leaves, because we assume that the size of I is at most the size of the queue. We call these leaves *target nodes*. The combiner then finds all *split nodes*: nodes for which the subtrees of both

children contain at least one target node. (See Figure 11 for an example of how target and split nodes can be defined.)



■ **Figure 11** Split and target nodes

sets $I[i].start$ to the right child of u_{i-1} (node $2 \cdot u_{i-1} + 1$) (Line 70). Then the combiner creates an InsertSet object X , sorts the arguments of the requests in I , puts them to $X.A$ and sets $a[1].split$ to X (Lines 81-89). Finally, it sets the status fields of all requests in I to SIFT (Lines 91-92).

Clients: Insert phase (Lines 41-85). Consider a client c with an incompleting request req . It waits while $a[req.start].split$ is null (Lines 42-43). Now c is going to insert values from InsertSet $a[req.start].split$ to the subtree of $req.start$. Let S be a local InsertSet variable initialized with $a[req.start].split$. For each node v on the path, c inserts values from S into the subtree of v . c calculates the minimum value x in S (Lines 51-53): the first element of $S.A$ or the first element of $S.B$. If $a[v].val$ is bigger than x , then the client removes x from S , appends $a[v].val$ to the end of $S.B$ and sets $a[v].val$ to x (Lines 54-60). Note that by the algorithm $S.B$ contains only values that were stored in the nodes above node v , thus, any value in $S.B$ cannot be bigger than $a[v].val$ and after appending $a[v].val$ $S.B$ remains sorted. Then the client calculates the number inL of target nodes in the subtree of the left child of v and the number inR of target nodes in the subtree of the right child of v (Lines 63-66, to calculate these numbers in constant time we use fields l and r of the request). If $inL = 0$, then all the values in S should be inserted into the subtree of the right child of v , and c proceeds with the right child $2v + 1$ (Lines 69-70). If $inR = 0$, then, symmetrically, c proceeds with the left child $2v$ (Lines 73-74). Otherwise, if $inL \neq 0$ and $inR \neq 0$, v is a split node and, thus, there is a client that waits at the right child $2v + 1$. Hence, c splits S to $(X, Y) \leftarrow S.split(inL)$ (Line 77): the values in X should be inserted into the subtree of node $2v$ and the values in Y should be inserted into the subtree of node $2v + 1$. Then c sets $a[2v + 1].split$ to Y , sets S to X and proceeds to node $2v$ (Lines 78-80). When c reaches a leaf v it sets the value $a[v].val$ to the only value in S (Lines 82-85) and sets the status of the request req to FINISHED (Line 7).

If the request of the combiner is an incompleting Insert, it runs the code above as a client (Figure 9 Lines 94-95). The combiner considers the Insert phase completed when all requests in I have status FINISHED (Lines 97-99).

Since we have $|I|$ target nodes, there are exactly $|I| - 1$ split nodes $u_1, \dots, u_{|I|-1}$: u_i is the lowest common ancestor of nodes $m + i$ and $m + i + 1$. They can be found in $O(|I| + \log m)$ primitive steps (Lines 55-72): starting with node $m + i$ go up the heap until a node becomes a left child of some node pr ; this pr is u_i . We omit the discussion about the fields l and r of $I[i]$: they represent the smallest and the largest leaf identifiers in the subtree of u_i , and they are used to calculate the number of leaves that are newly inserted, i.e., $m + 1, \dots, m + |I|$, in a subtree in constant time. The combiner sets $I[1].start$ to the root (the node with identifier 1), (Line 55) and, for each $i \in [2, |I|]$, it

A.6 Analysis

Now we provide theorems on correctness and time bounds.

► **Theorem 2.** *Our concurrent priority queue implementation is linearizable.*

Proof. The execution can be split into combining phases (Figure 9 Lines 1-99) which do not intersect. We group the operations into batches corresponding to the combining phase in which they are applied.

Consider the i -th combining phase. We linearize all the operations from the i -th phase right after the end of the corresponding `getRequests()` (Line 1) in the following order: at first, we linearize `ExtractMin` operations in the increasing order of their responses, then, we linearize `Insert` operations in any order.

To see that this linearization is correct it is enough to prove that the combiner and the clients apply the batch correctly.

► **Lemma 3.** *Suppose that the batch of the i -th combining phase contains a `ExtractMin` operations and b `Insert` operations with arguments x_1, \dots, x_b . Let V be the set of values stored in the priority queue before the i -th phase. The combiner and the clients apply this batch correctly:*

- *The minimal a values y_1, \dots, y_a in V are returned to `ExtractMin` operations.*
- *After an execution the set of values stored in the queue is equal to $V \cup \{x_1, \dots, x_b\} \setminus \{y_1, \dots, y_a\}$. and the values are stored in nodes with identifiers $1, \dots, |V| - b + a$.*
- *After an execution the heap property is satisfied for each node.*

Proof. The first statement is correct, because the combiner chooses the smallest a elements from the priority queue and sets them as the results of `ExtractMin` requests (Lines 18-33).

The second statement about the set of values straightforwardly follows from the algorithm. During `ExtractMin` phase the combiner finds a smallest elements, replaces them with $x_1, \dots, x_{\min(a,b)}$ and with values from the last $a - \min(a,b)$ nodes of the heap: the set of values in the priority queue becomes $V \cup \{x_1, \dots, x_{\min(a,b)}\} \setminus \{y_1, \dots, y_b\}$ and the values are stored in nodes $1, \dots, |V| - a + \min(a,b)$. Then, the *sift down* is initiated, but it does not change the set of values and it does not touch nodes other than $1, \dots, |V| - a + \min(a,b)$. During `Insert` phase the values $x_{\min(a,b)+1}, \dots, x_b$ are inserted and new nodes which are used in `Insert` phase are $|V| - a + \min(a,b) + 1, \dots, |V| - a + b$. Thus, the final set of values is $V \cup \{x_1, \dots, x_b\} \setminus \{y_1, \dots, y_a\}$ and the values are stored in nodes $1, \dots, |V| - a + b$.

The third statement is slightly tricky. At first, the combiner finds a smallest elements that should be removed and replaces them with $x_1, \dots, x_{\min(a,b)}$ and with values from the last $a - \min(a,b)$ nodes of the heap. Suppose that these a smallest elements were at nodes v_1, \dots, v_a , sorted by their depth (the length of the shortest path from the root) in non-increasing order. These nodes form a connected subtree where v_a is the root of the heap. Suppose that they do not form a connected tree or v_a is not the root of the heap. Then there exists a node v_i which parent p is not v_j for any j . This means that the values in nodes v_1, \dots, v_a are not the smallest a values: by the heap property a value in p is smaller than the value in v_i .

Now a processes perform *sift down* from the nodes v_1, \dots, v_a . We show that when a node v is unlocked, i.e., its `locked` field is set to false, the value at v is the smallest value in the subtree of v . This statement is enough to show that the heap property holds for all nodes after `ExtractMin` phase, because at that point all nodes are unlocked.

Consider an execution of *sift down*. We prove the statement by induction on the number of unlock operations. Base. No unlock happened and the statement is satisfied for all

unlocked nodes, i.e., all the nodes except for v_1, \dots, v_a . Transition. Let us look right before the k -th unlock: the unlock of a node v . The left child l of v should be unlocked and, thus, l contains a value that is the smallest in its subtree. The same statement holds for the right child r of v . v chooses the smallest value between the value at v and the values at l and r . This value is the smallest in the subtree of v . Thus, the statement holds for v when unlocked.

After that, the algorithm applies the incompleting $b - \min(a, b)$ Insert operations. We name the nodes with at least one target node in the subtree as *modified*. Modified nodes are the only nodes whose value can be changed and, also, each modified node is visited by exactly one client. To prove that after the execution the heap property for each modified node holds: we show by induction on the depth of a modified node that if a node v is visited by a client with InsertSet S then: (1) $S.A$ is sorted; (2) $S.B$ is sorted and contains only values that were stored in ancestors of v after ExtractMin phase; and (3) v contains the smallest value in its subtree when the client finishes with it. Base. In the root $S.A$ is sorted, $S.B$ is empty and the new value in the root is either the first value in $S.A$ or the current value in the root, thus, it is the smallest value in the heap. Transition from depth k to depth $k + 1$. Consider a modified node v at depth $k + 1$ and its parent p . Suppose that p was visited by a client with InsertSet S_p . By induction, $S_p.A$ is sorted and $S_p.B$ is sorted and contains only the values that were in ancestors. Then the client chooses the smallest value in p : either $a[p]$, the first value of $S_p.A$ or the first value of $S_p.B$. Note that after any of these three cases $S_p.A$ and $S_p.B$ are sorted and $S_p.B$ contains only values from ancestors and node p :

- $a[p]$ is the smallest, then $S_p.A$ and $S_p.B$ are not modified;
- we poll the first element of $S_p.A$ or $S_p.B$; $S_p.A$ and $S_p.B$ are still sorted; then we append $a[p]$ to $S_p.B$, and $a[p]$ has to be the biggest element in $S_p.B$, since $S_p.B$ contains only the values from ancestors.

Then the client splits S_p and some client, possibly, another one, works on v with IntegerSet S . Since, S is a subset of S_p then $S.A$ is sorted and $S.B$ is sorted and contains only the values from ancestors (ancestors of p and, possibly, p). Finally, the client chooses the smallest value to appear in the subtree: the first value of $S.A$, the first value of $S.B$ and $a[v]$. ◀

◀

► **Theorem 4.** *Suppose that the combiner collects c requests using `getRequests()`. Then the combiner and the clients apply these requests to a priority queue of size m using $O(c + \log m)$ RMRs in CC model each and $O(c \cdot (\log c + \log m))$ RMRs in CC model in total.*

Proof. Suppose that the batch consists of a ExtractMin operations and b Insert operations.

The combiner splits requests into two sets E and I ($O(c)$ RMRs, Lines 9-16). Then it finds a nodes with the smallest values ($O(a \log a)$ primitive steps, but $O(a)$ RMRs, Lines 18-28) using Dijkstra-like algorithm. After that, the combiner sets up ExtractMin requests, sets their status to SIFT and pairs some Insert requests with ExtractMin requests ($O(a)$ RMRs, Lines 30-45).

The clients participate in ExtractMin phase. At first, each client waits for its status to change (1 RMR). Then the client performs at most $\log m$ iterations of the loop (Line 12): waits on the *locked* fields of the children ($O(1)$ RMRs, Lines 13-17); reads the values in the children ($O(1)$ RMRs, Line 20); compares these values with the value at the node, possibly, swap the values, lock the proper child and unlock the node ($O(1)$ RMRs, Lines 22-29). When the client stops it changes the status (1 RMR, Line 7).

The combiner waits for the change of the status of the clients ($O(a)$ RMRs, Lines 50-52). Summing up, in ExtractMin phase each client performs $O(\log m)$ RMRs and the combiner performs $O(a + \log m)$ RMRs, giving $O(c + c \cdot \log m)$ RMRs in total.

The combiner throws away completed Insert requests ($O(b)$ primitive steps and 0 RMRs, Line 53). Then it finds the split nodes ($O(\log m + b)$ primitive steps, but 0 RMRs, Lines 55-72). After that the combiner sorts arguments of remaining Insert requests, sets their status to SIFT and sets up the initial InsertSet ($O(b \cdot \log b)$ primitive steps, but $O(b)$ RMRs, Line 81 and Line 92).

The clients participate in Insert phase. At first, a client t waits while the corresponding InsertSet is null (1 RMR, Lines 41-43). Suppose that it reads the InsertSet S and starts the traversal down. The client performs at most $\log m$ iterations of the loop (Line 50): choose the smallest value ($O(1)$ RMRs, Lines 51-60), find whether to split InsertSet ($O(1)$ RMRs, Lines 62-74), split InsertSet (calculated below, Line 77) and pass one InsertSet to another client ($O(1)$ RMRs, Lines 78-80). Now let us calculate the number of RMRs spent in Line 77. Suppose that there are k iterations of the loop and the size of S at iteration i is s_i . At the i -th iteration split works in $O(\min(s_{i+1}, s_i - s_{i+1})) = O(s_i - s_{i+1})$ primitive steps and RMRs. Summing up through all iterations we get $O(s_1) = O(b)$ RMRs spent by t in Line 77. Finally, t sets the value in the leaf ($O(1)$ RMRs, Lines 82-85) and changes the status ($O(1)$ RMRs, Line 7).

The combiner waits for the change of the status of the clients ($O(b)$ RMRs, Lines 97-99). Summing up, in Insert phase the clients and the combiner perform $O(b + \log m)$ RMRs each. Consequently, the straightforward bound on the total number of RMRs is $O(c^2 + c \cdot \log m)$ RMRs.

To get the improved bound we carefully calculate the total number of RMRs spent on the splits of InsertSets in Line 77. This number equals to the number of values that are moved to newly created sets during the splits. For simplicity, we assume that inserted values are bigger than all the values in the priority queue and, thus, each InsertSet contains only the newly inserted values. This assumption does not affect the bound. Consider now the inserted value v . Suppose that v was moved k times and at the i -th time it was moved during the split of InsertSet with size s_i . Because v is moved during split only to the set with the smaller size: $s_1 \geq 2 \cdot s_2 \geq \dots \geq 2^{k-1} \cdot s_k$. k is less than $\log c$, because $s_1 \leq c$, and, thus, v was moved no more than $\log c$ times. This means, that in total during the splits of InsertSets no more than $c \cdot \log c$ values are moved to new sets, giving $O(c \cdot \log c)$ RMRs during the splits. This gives us a total bound of $O(c \cdot (\log c + \log m))$ RMRs during Insert phase.

To summarize, the combiner and the clients perform $O(c + \log m)$ RMRs each and $O(c \cdot (\log c + \log m))$ RMRs in total. ◀

► **Remark.** The above bounds also hold in DSM model for the version of the described algorithm. For that we have to simply make spin-loops to loop on the local variable of processes. In our algorithm the purpose of each spin-loops is to wake up some process. At most places in our algorithm when we set the variable on which we spin we know (or can deduce by a simple modification of the algorithm) which process is going to wake up. For each spin-loop it is enough to create a separate variable in the memory of the target process.

The only two non-trivial spin-loops are in CLIENT_CODE (Lines 13-17) where we do not know a process that is going to wake up. To obviate this issue we expand each Node object with the pointer to process $proc$. When the thread wants to *sift-down*, first, it registers itself in $a[v].proc$ and, then, checks $a[2v].locked$ and $a[2v+1].locked$. If some of them are **true** then it spins on specifically created local variables: on $notify_{2v}$ if $a[2v].locked$ is **true**, and

0:26 Parallel Combining

on $notify_{2v+1}$ if $a[2v+1].locked$ is **true**. Then, the algorithm standardly performs swapping routine. At the end, it unlocks the node, i.e., sets $a[v].locked$ to **false**, then, reads a process $a[v/2].proc$ and notifies it by setting its corresponding variable $notify_v$. Note that the total number of *notify* local variables that is needed by each process is logarithmic from the size of the queue.

The described transformation (in reality, it is slightly more technical than described above) of our algorithm provides an algorithm with the same bounds on RMRs but in DSM model.