



Decentralized Online Scheduling of Malleable NP-hard Jobs

Peter Sanders^(✉)  and Dominik Schreiber 

Karlsruhe Institute of Technology, 76131 Karlsruhe,
Germany

{sanders,dominik.schreiber}@kit.edu



Abstract. In this work, we address an online job scheduling problem in a large distributed computing environment. Each job has a priority and a demand of resources, takes an unknown amount of time, and is malleable, i.e., the number of allotted workers can fluctuate during its execution. We subdivide the problem into (a) determining a fair amount of resources for each job and (b) assigning each job to an according number of processing elements. Our approach is fully decentralized, uses lightweight communication, and arranges each job as a binary tree of workers which can grow and shrink as necessary. Using the NP-complete problem of propositional satisfiability (SAT) as a case study, we experimentally show on up to 128 machines (6144 cores) that our approach leads to near-optimal utilization, imposes minimal computational overhead, and performs fair scheduling of incoming jobs within a few milliseconds.

Keywords: Malleable job scheduling · Load balancing · SAT

1 Introduction

A parallel task is called *malleable* if it can handle a fluctuating number of workers during its execution. In the field of distributed computing, malleability has long been recognized as a powerful paradigm which opens up vast possibilities for fair and flexible scheduling and load balancing [13, 17]. While most previous research on malleable job scheduling has steered towards iterative data-driven applications, we want to shed light on malleability in a very different context, namely for NP-hard tasks with unknown processing times. For instance, the problem of propositional satisfiability (SAT) is of high practical relevance and an important building block for many applications including automated planning [26], formal verification [18], and cryptography [20]. We consider malleable scheduling of such tasks highly promising: On the one hand, the description of a job can be relatively small even for very difficult problems, and the successful approach of employing many combinatorial search strategies in parallel can be made malleable without redistribution of data [27]. On the other hand, the limited scalability of these parallel algorithms calls for careful distribution of computational resources. We believe that a cloud-like on-demand system for

resolving NP-hard problems has the potential to drastically improve efficiency and productivity for many organizations and environments. Using malleable job scheduling, we can schedule new jobs within a few milliseconds, resolve trivial jobs in a fraction of second, and rapidly resize more difficult jobs to a fair share of all resources – as far as the job can make efficient use of these resources.

To meet these objectives, we propose a fully decentralized scheduling approach which guarantees fast, fair, and bottleneck-free scheduling of resources without any knowledge on processing times. In previous work [27], we briefly outlined initial algorithms for this purpose while focusing on our award-winning scalable SAT solving engine which we embedded into our system. In this work, we shed more light on our earlier scheduling algorithms and proceed to propose significant improvements both in theory and in practice.

We address two subproblems. The first problem is to let m workers compute a fair number of workers v_j for each active job j , accounting for its priority and maximum demand, which result in optimal system utilization. In previous work [27] we outlined this problem and employed a black box algorithm to solve it. The second problem is to assign v_j workers to each job j while keeping the assignment as stable as possible over time. Previously [27], we proposed to arrange each job j as a binary tree of workers which grows and shrinks depending on v_j , and we described and implemented a worker assignment strategy which routes request messages randomly through the system. When aiming for optimal utilization, this protocol leads to high worst-case scheduling latencies.

In this work, we describe fully distributed and bottleneck-free algorithms for both of the above problems. Our algorithms have $\mathcal{O}(\log m)$ span and are designed to consistently achieve optimal utilization. Furthermore, we introduce new measures to preferably reuse existing (suspended) workers for a certain job rather than initializing new workers. We then present our scheduling platform *Mallob*¹ which features simplified yet highly practical implementations of our approaches. Experiments on up to 128 nodes (6144 cores) show that our system leads to near-optimal utilization and schedules jobs with a fair share of resources within tens of milliseconds. We consider our theoretical as well as practical results to be promising contributions towards processing malleable NP-hard tasks in a more scalable and resource-efficient manner.

2 Preliminaries

We now establish important preliminaries and discuss work related to ours.

2.1 Malleable Job Scheduling

We use the following definitions [10]: A *rigid* task requires a fixed number of workers. A *moldable* task can be scaled to a number of workers at the time of its scheduling but then remains rigid. Finally, a *malleable* task is able to adapt to a fluctuating number of workers *during* its execution. Malleability can be a

¹ <https://github.com/domschrei/mallob>.

highly desirable property of tasks because it allows to balance tasks continuously to warrant fair and optimal utilization of the system at hand [17]. For instance, if an easy job arrives in a fully utilized system, malleable scheduling allows to shrink an active job in order to schedule the new job immediately, significantly decreasing its response time. Due to the appeal of malleable job scheduling, there has been ongoing research to exploit malleability, from shared-memory systems [13] to HPC environments [6, 9], even to improve energy efficiency [25].

The effort required to transform a moldable (or rigid) algorithm into a malleable algorithm depends on the application at hand. For iterative data-driven applications, redistribution of data is necessary if a task is expanded or shrunk [9]. In contrast, we demonstrated in previous work [27] for the use case of propositional satisfiability (SAT) that basic malleability is simple to achieve if the parallel algorithm is composed of many independent search strategies: The abrupt suspension and/or termination of individual workers can imply the loss of progress, but preserves completeness. Moreover, if workers periodically exchange knowledge, the progress made on a worker can benefit the job even if the worker is removed. For these reasons, we have not yet considered the full migration of application processes as is done in adaptive middlewares [9, 16] but instead hold the application itself responsible to react to workers being added or removed.

Most prior approaches rely on known processing times of jobs and on an accurate model for their execution time relative to the degree of parallelism [5, 24] whereas we do not rely on such knowledge. Furthermore, most approaches employ a centralized scheduler, which implies a potential bottleneck and a single point of failure. Our approach is fully decentralized and uses a small part of each processes' CPU time to perform distributed scheduling, which also opens up the possibility to add more general fault-tolerance to our work in the future. For instance, this may include continuing to schedule and process jobs correctly even in case of network-partitioning faults [2], i.e., failures where sub-networks in the distributed environment are disconnected from each another. Other important aspects of fault-tolerance include mitigation of simple node failures (i.e., a machine suddenly goes out of service) and of Byzantine failures [7] (i.e., a machine exhibits arbitrary behavior, potentially due to a malicious attack).

2.2 Scalable SAT Solving

The propositional satisfiability (SAT) problem poses the question whether a given propositional formula $F = \bigwedge_{i=1}^k (\bigvee_{j=1}^{c_i} l_{i,j})$ is satisfiable, i.e., whether there is an assignment to all Boolean variables in F such that F evaluates to **true**. SAT is the archetypical NP-complete problem [8] and, as such, a notoriously difficult problem to solve. SAT solving is a crucial building block for a plethora of applications such as automated planning [26], formal verification [18], and cryptography [20]. State-of-the-art SAT solvers are highly optimized: The most popular algorithm named Conflict-Driven Clause Learning (CDCL) performs depth-first search on the space of possible assignments, backtracks and restarts its search frequently, and derives redundant *conflict clauses* when encountering a dead end in its search [19]. As these clauses prune search space

and can help to derive unsatisfiability, remembering important derived clauses is crucial for modern SAT solvers' performance [3].

The empirically best performing approach to parallel SAT solving is a so-called *portfolio* of different solver configurations [14] which all work on the original problem and periodically exchange learned clauses. In previous work, we presented a highly competitive portfolio solver with clause sharing [27] and demonstrated that careful periodic clause sharing can lead to respectable speedups for thousands of cores. The malleable environment of this solver is the system which we present here. Other recent works on decentralized SAT solving [15, 21] rely on a different parallelization which generates many independent subproblems and tends to be outperformed by parallel portfolios for most practical inputs [11].

2.3 Problem Statement

We consider a homogeneous computing environment with a number of interconnected machines on which a total of m *processing elements*, or PEs in short, are distributed. Each PE has a *rank* $x \in \{0, \dots, m-1\}$ and runs exclusively on $c \geq 1$ cores of its local machine. PEs can only communicate via message passing.

Jobs are introduced over an interface connecting to some of the PEs. Each job j has a job description, a priority $p_j \in \mathbb{R}^+$, a *demand* $d_j \in \mathbb{N}^+$, and a *budget* b_j (in terms of wallclock time or CPU time). If a PE participates in processing a job j , it runs an execution environment of j named a *worker*. A job's demand d_j indicates the maximum number of parallel workers it can currently employ: d_j is initialized to 1 and can then be adjusted by the job after an initial worker has been scheduled. A job's priority p_j may be set, e.g., depending on who submitted j and on how important they deem j relative to an average job of theirs. In a simple setting where all jobs are equally important, assume $p_j = 1 \forall j$. A job is cancelled if it spends its budget b_j before finishing. We assume for the active jobs J in the system that the number $n = |J|$ of active jobs is no higher than m and that each PE employs at most one worker at any given time. However, a PE can preempt its current worker, run a worker of another job, and possibly resume the former worker at a later point.

Let T_j be the set of active workers of $j \in J$. We call $v_j := |T_j|$ the *volume* of j . Our aim is to continuously assign each $j \in J$ to a set T_j of PEs subject to:

- (C1) (Optimal utilization) Either all job demands are fully met or all m PEs are working on a job: $(\forall j \in J : v_j = d_j) \vee \sum_{j \in J} v_j = m$.
- (C2) (Individual job constraints) Each job must have at least one worker and is limited to d_j workers: $\forall j \in J : 1 \leq v_j \leq d_j$.
- (C3) (Fairness) Resources allotted to each job j scale proportionally with p_j except if prevented by C2:
For each $j, j' \in J$ with $p_j \geq p_{j'}$, there are *fair assignments* $\omega, \omega' \in \mathbb{R}^+$ with $\omega/\omega' = p_j/p_{j'}$ and some $0 \leq \varepsilon \leq 1$ such that $v_j = \min(d_j, \max(1, \lfloor \omega + \varepsilon \rfloor))$ and $v_{j'} = \min(d_{j'}, \max(1, \lfloor \omega' \rfloor))$.

Due to rounding, in C3 we allow for job volumes to deviate by a single unit (see $\varepsilon \leq 1$) from a fair distribution as long as the job of higher priority is favored.

3 Approach

We subdivide the problem at hand into two subproblems: First, find *fair volumes* v_j for all currently active jobs $j \in J$ subject to C1–C3. Secondly, identify pairwise disjoint sets T_j with $|T_j| = v_j$ for each $j \in J$. In this section, we present fully decentralized and highly scalable algorithms for both subproblems. In Sect. 4.1 we describe how our practical implementation differs from these algorithms.

To assess our algorithms, we consider two important measures from parallel processing. Given a distributed algorithm, consider the dependency graph which is induced by the necessary communication among all PEs. The *span* (or *depth*) of the algorithm is the length of a critical path through this graph. The *local work* is the complexity of local computations summed up over all PEs.

3.1 Calculation of Fair Volumes

Given jobs J with individual priorities and demands, we want to find a fair volume v_j for each job j such that constraints C1–C3 are met. Volumes are recomputed periodically taking into account new jobs, departing jobs, and changed demands. In the following, assume that each job has a single worker which represents this (and only this) job. We elaborate on these representants in Sect. 3.2.

We defined our problem such that $n = |J| \leq m$. Similarly, we assume $\sum_{j \in J} d_j > m$ since otherwise we can trivially set $v_j = d_j$ for all jobs j . Assuming real-valued job volumes for now, we can observe that for any parameter $\alpha \geq 0$, constraints C2–C3 are fulfilled if we set $v_j = v_j(\alpha) := \max(1, \min(d_j, \alpha p_j))$. By appropriately choosing α , we can also meet the utilization constraint C1: Consider the function $\xi(\alpha) := m - \sum_{j \in J} v_j(\alpha)$ which expresses the unused resources for a particular value of α . Function ξ is a continuous, monotonically decreasing, and piece-wise linear function (see Fig. 1). Moreover, $\xi(0) = m - n \geq 0$ and




Fig. 1. Volume calculation example with four jobs and $m = 7$. Five of the eight points where $\xi(\alpha)$ is evaluated are depicted, three more (d_3/p_3 , d_1/p_1 , and d_2/p_2) are omitted. In the interval $[1/p_2, 1/p_1]$ we find $\alpha_0 = 0.8$ (red circle) where $\xi(\alpha) = 0$. Job 4 is capped at its demand ($v_4 = 2$) and job 1 is raised to $v_1 = 1$. The real-valued shares $\alpha_0 p_2 = 1.6$ and $\alpha_0 p_3 = 2.4$ are rounded to $v_2 = 1$ and $v_3 = 3$ as job 3 has the higher priority. (Color figure online)

$\xi(\max_{j \in J} d_j/p_j) = m - \sum_{j \in J} d_j < 0$. Hence $\xi(\alpha) = 0$ has a solution α_0 which represents the desired choice of α that exploits all resources, i.e., it also fulfills constraint C1. Once α_0 is found, we need to round each $v_j(\alpha_0)$ to an integer. Due to C1 and C3, we propose to round down all volumes and then increment the volume of the $k := m - \sum_j \lfloor v_j(\alpha_0) \rfloor$ jobs of highest priority: We identify $J' := \{j \in J : v_j(\alpha_0) < d_j\}$, sort J' by job priority, and select the first k jobs.

We now outline a fully distributed algorithm which finds α_0 in logarithmic span. We exploit that ξ' , the gradient of ξ , changes at no more than $2n$ values of α , namely when $\alpha p_j = 1$ or $\alpha p_j = d_j$ for some $j \in J$. Since we have $m \geq n$ PEs available, we can try these $\mathcal{O}(n)$ values of $\xi(\alpha)$ in parallel. We then find the two points with smallest positive value and largest negative value using a parallel reduction operation. Lastly, we interpolate ξ between these points to find α_0 .

The parallel evaluation of ξ is still nontrivial since a naive implementation would incur quadratic work – $\mathcal{O}(n)$ for each value of α . We now explain how to accelerate the evaluation of ξ . For this, we rewrite $\xi(\alpha) = m - \sum_{j \in J} v_j(\alpha)$ as:

$$\xi(\alpha) = m - \underbrace{\left(\sum_{j: \alpha p_j < 1} 1 + \sum_{j: \alpha p_j > d_j} d_j \right)}_R - \alpha \underbrace{\sum_{j: 1 \leq \alpha p_j \leq d_j} p_j}_P \quad (1)$$

Intuitively, R sums up all resources which are assigned due to raising a job volume to 1 (if $\alpha p_j < 1$) and due to capping a job volume at d_j (if $\alpha p_j > d_j$); and αP sums up all resources assigned as $v_j = \alpha p_j$ (if $1 \leq \alpha p_j \leq d_j$).

This new representation only features two unknown variables, R and P , which can be computed efficiently. At $\alpha = 0$, we have $R = n$ and $P = 0$ since all job volumes are raised to one. If we then successively increase α , we pass $2n$ events where R and P are modified, namely whenever $\alpha p_j = 1$ or $\alpha p_j = d_j$ for some job j . Since each such event modifies R and P by a fixed amount, we can use a single prefix sum calculation to obtain *all* intermediate values of R and P .

Each event $e = (\alpha_e, r_e, p_e)$ occurs at point α_e and adds r_e to R and p_e to P . Each job j causes two events: $\underline{e}_j = (1/p_j, -1, p_j)$ for the point $\alpha p_j = 1$ where v_j stops being raised to 1, and $\bar{e}_j = (d_j/p_j, d_j, -p_j)$ for the point $\alpha p_j = d_j$ where v_j begins to be capped at d_j . We sort all events by α_e and then compute a prefix sum over r_e and p_e : $(R_e, P_e) = (\sum_{e' \preceq e} r_{e'}, \sum_{e' \preceq e} p_{e'})$, where “ \preceq ” denotes the ordering of events after sorting. We can now compute $\xi(\alpha_e) = m - (n + R_e) - \alpha_e P_e$ at each event e .² The value of n can be obtained with a parallel reduction.

Overall, our algorithm has $\mathcal{O}(\log m)$ span and takes $\mathcal{O}(m \log m)$ work: Sorting $\mathcal{O}(n)$ elements in parallel on $m \geq n$ PEs is possible in logarithmic time,³ as is computing reductions and prefix sums. Selecting the k jobs to receive additional volume after rounding down all volumes can be reduced to sorting as well.

² If there are multiple events at the same α , their prefix sum results can differ but will still result in the same $\xi(\alpha)$. This is due to the continuous nature of ξ : Note how each event modifies the gradient $\xi'(\alpha)$ but preserves the value of $\xi(\alpha)$.

³ Asymptotically optimal sorting on communication networks [1] is of mostly theoretical value due to the large constant values involved. However there are quite practical algorithms when $n \in \mathcal{O}(\sqrt{m})$ or when spending $\mathcal{O}(\log^2 n)$ time is acceptable [4].




Fig. 2. Left: Job tree T_j features ten workers $\{w_j^0, w_j^1, \dots, w_j^9\}$ due to the volume $v_j = 10$ assigned to j . Right: Volume update $v_j = 4$ arrives. Consequently, all workers with index ≥ 4 are suspended and the corresponding PEs can adopt another job.

3.2 Assignment of Jobs to PEs

We now describe how the fair volumes computed as in the previous section translate to an actual assignment of jobs to PEs.

Basic Approach. We begin with our basic approach as introduced in [27].

For each job j , we address the k current workers in T_j as $w_j^0, w_j^1, \dots, w_j^{k-1}$. These workers can be scattered throughout the system, i.e., their *job indices* $0, \dots, k-1$ within T_j are not to be confused with their ranks. The k workers form a communication structure in the shape of a binary tree (Fig. 2). Worker w_j^0 is the root of this tree and represents j for the calculation of its volume (Sect. 3.1). Workers w_j^{2i+1} and w_j^{2i+2} are the left and right children of w_j^i . Jobs are made malleable by letting T_j grow and shrink dynamically. Specifically, we enforce that T_j consists of exactly $k = v_j$ workers. If v_j is updated, all workers w_j^i for which $i \geq v_j$ are suspended and the corresponding PEs turn idle. Likewise, workers without a left (right) child for which $2i+1 < v_j$ ($2i+2 < v_j$) attempt to find a child worker w_j^{2i+1} (w_j^{2i+2}). New workers are found via *request messages*: A request message $r = (j, i, x)$ holds index i of the requested worker w_j^i as well as rank x of the requesting worker. If a new job is introduced at some PE, then this PE emits a request for the root node w_j^0 of T_j . All requests for w_j^i , $i > 0$ are emitted by the designated parent node $w_j^{\lfloor (i-1)/2 \rfloor}$ of the desired worker.

In [27], we proposed that each request performs a random walk through a regular graph of all PEs and is resolved as soon as it hits an idle PE. While this strategy resolves most requests quickly, some requests can require a large number of hops. If we assume a fully connected graph of PEs and a small share ϵ of workers is idle, then each hop of a request corresponds to a Bernoulli process with success probability ϵ , and a request takes an expected $1/\epsilon$ hops until an idle PE is hit. Consequently, to improve worst-case latencies, a small ratio of workers should be kept idle [27]. By contrast, our following algorithm with logarithmic span does not depend on suboptimal utilization.

Matching Requests and Idle PEs. In a first phase, our improved algorithm (see Fig. 3) computes two prefix sums with one collective operation: the number




Fig. 3. Examples for matching requests and idle PEs. White (gray) squares represent idle (busy) PEs, spheres (diamonds) represent requests (idle tokens). Left: A prefix sum (not depicted) numbers all requests and idle tokens, and each request (token) of index i is sent to rank i . Each PE with a matching pair sends the request to the idle PE. Right: A job j grows by multiple layers of T_j . Requests are sent along a tree structure and child-parent relationships of T_j are encoded into the distributed requests.

of requests q_i being emitted by PEs of rank $< i$, and the number o_i of idle PEs of rank $< i$. We also compute the total sums, q_m and o_m , and communicate them to all PEs. The q_i and o_i provide an implicit global numbering of all requests and all idle PEs. In a second phase, the i -th request and the i -th token are both sent to rank i . In the third and final phase, each PE which received both a request and an idle token sends the request to the idle PE referenced by the token.

If the request for a worker w_j^i is only emitted by its designated parent $w_j^{[(i-1)/2]}$, then our algorithm so far may need to be repeated $\mathcal{O}(\log m)$ times: Repetition l activates a worker which then emits requests for repetition $l + 1$. Instead, we can let a worker emit requests not only for its direct children, but for all *transitive* children it deserves. Each worker w_j^i can compute the number k of desired transitive children from v_j and i . The worker then contributes k to q_i . In the second phase, the k requests can be distributed communication-efficiently to a range of ranks $\{x, \dots, x + k - 1\}$: w_j^i sends requests for workers w_j^{2i+1} and w_j^{2i+2} to ranks x and $x + 1$, which send requests for corresponding child workers to ranks $x + 2$ through $x + 5$, and so on, until worker index $v_j - 1$ is reached. To enable this distribution, we append to each request the values x , v_j , and the rank of the PE where the respective parent worker will be initialized. As such, each child knows its parent within T_j (Fig. 3) for job-internal communication.

We now outline how our algorithm can be executed in a fully asynchronous manner. We compute the prefix sums within an In-Order binary tree of PEs [22, Chapter 13.3], that is, all children in the left subtree of rank i have a rank $< i$ and all children in the right subtree have a rank $> i$. This prefix sum computation can be made sparse and asynchronous: Only non-zero contributions to a prefix sum are sent upwards explicitly, and there is a minimum delay in between sending contributions to a parent. Furthermore, we extend our prefix sums to also include *inclusive* prefix sums q'_i, o'_i which denote the number of requests (tokens) at PEs of rank $\leq i$. As such, every PE can see from the difference $q'_i - q_i$ ($o'_i - o_i$) how many of its local requests (tokens) took part in the prefix sum. Last but not

least, the number of tokens and the number of requests may not always match – a PE which receives either a request or an idle token (but not both) knows of this imbalance due to the total sums q_m, o_m . The unmatched message is sent to its origin and can re-participate in the next iteration.

Our matching algorithm has $\mathcal{O}(\log m)$ span and takes $\mathcal{O}(m)$ local work. The maximum local work of any given PE is in $\mathcal{O}(\log m)$ (to compute the above k), which is amortized by other PEs because at most m requests are emitted.

3.3 Reuse of Suspended Workers

Each PE remembers up to C most recently used workers (for a small constant C) and deletes older workers. Therefore, if a worker w_j^i is suspended, it may be resumed at a later time. Our algorithms so far may choose different PEs and hence create new workers whenever T_j shrinks and then re-grows. We now outline how we can increase the reuse of suspended workers.

In our previous approach [27], each worker remembers a limited number of ranks of its past (direct) children. A worker which desires a child queries them for reactivation one after the other until success or until all past children have been queried unsuccessfully, at which point a normal job request is emitted.

We make two improvements to this strategy. First, we remember past workers in a distributed fashion. More precisely, whenever a worker joins or leaves T_j , we distribute information along T_j to maintain the following invariant: Each current leaf w_j^i in T_j remembers the past workers which were located in a subtree below index i . As such, past workers can be remembered and reused even if T_j shrinks by multiple layers and re-grows differently.

Secondly, we adjust our scheduling to actively prioritize the reuse of existing workers over the initialization of new workers. In our implementation, each idle PE can infer from its local volume calculation (Sect. 4.1) which of its local suspended workers w_j^i are eligible for reuse, i.e., $v_j > i$ in the current volume assignment. If a PE has such a worker w_j^i , the PE will reject any job requests until it received a message regarding w_j^i . This message is either a query to resume w_j^i or a notification that w_j^i will not be reused. On the opposite side, a worker which desires a child begins to query past children according to a “most recently used” strategy. If a query succeeds, all remaining past children are notified that they will not be reused. If all queries failed, a normal job request is emitted.

4 The Mallob System

In the following, we outline the design and implementation of our platform named Mallob, short for **M**alleable **L**oad **B**alancer. Mallob is a C++ application using the Message Passing Interface (MPI) [12]. Each PE can be configured to accept jobs and return responses, e.g., over the local file system or via an internal API. The application-specific worker running on each PE is defined via an interface with a small set of methods. These methods define the worker’s behavior if it is started, suspended, resumed, or terminated, and allow it to send and receive

application-specific messages at will. Note that we outlined some of Mallob’s earlier features in previous work [27] with a focus on our malleable SAT engine.

4.1 Implementation of Algorithms

Our system features practical and simplified implementations solving the volume assignment problem and the request matching problem. We now explain how and why these implementations differ from the algorithms provided in Sect. 3.

Volume Assignment. Our implementation computes job volumes similar to the algorithm outlined in Sect. 3.1. However, each PE computes the desired change of root α_0 of ξ locally. All events in the system (job arrivals, departures, and changes in demands) are aggregated and broadcast periodically such that each PE can maintain a local image of all active jobs’ demands and priorities [27]. The local search for α_0 is then done via bisection over the domain of ξ . This approach requires more local work than our fully distributed algorithm and features a broadcast of worst-case message length $\mathcal{O}(n)$. However, it only requires a single all-reduction. At the scale of our current implementation ($n < 10^3$ and $m < 10^4$), we expect that our simplified approach performs better than our asymptotically superior algorithm which features several stages of collective operations. When targeting much larger configurations in the future, it may be beneficial to implement and employ our fully distributed algorithm instead.

Request Matching. We did not yet implement asynchronous prefix sums as described in Sect. 3.2. Instead, we route requests directly along a communication tree R of PEs. Each PE keeps track of the *idle count*, i.e., the number of idle PEs, in each of its subtrees in R . This count is updated transitively whenever the idle status of a child changes. Emitted requests are routed upwards through R until hitting an idle PE or until a hit PE has a subtree with a non-zero idle count, at which point the request is routed down towards the idle PE. If a large number of requests (close to n) are emitted, the traffic along the root of R may constitute a bottleneck. However, we found that individual volume updates in the system typically result in a much smaller number of requests, hence we did not observe such a bottleneck in practice. We intend to include our bottleneck-free algorithm (Sect. 3.2) in a future version of our system.

4.2 Engineering

For good practical performance of our system, careful engineering was necessary. For instance, our system exclusively features asynchronous communication, i.e., a PE will never block for an I/O event when sending or receiving messages. As a result, our protocols are designed without explicit synchronization (barriers or similar). We only let the main thread of a PE issue MPI calls, which is the most widely supported mode of operation for multithreaded MPI programs.

As we aim for scheduling latencies in the range of milliseconds, each PE must frequently check its message queue and react to messages. For instance, if the main thread of a PE allocates space for a large job description, this can cause a prohibitively long period where no messages are processed. For this reason, we use a separate thread pool for all tasks which involve a risk of taking a long time. Furthermore, we split large messages into batches of smaller messages, e.g., when transferring large job descriptions to new workers.

5 Evaluation

We now present our experimental evaluation. All experiments have been conducted on the supercomputer SuperMUC-NG. If not specified otherwise, we used 128 compute nodes, each with an Intel Skylake Xeon Platinum 8174 processor clocked at 2.7 GHz with 48 physical cores (96 hardware threads) and 96 GB of main memory. SuperMUC-NG is running Linux (SLES) with kernel version 4.12 at the time of running our experiments. We compiled Mallob with GCC 9 and with Intel MPI 2019. We launch twelve PEs per machine, assign eight hardware threads to each PE, and let a worker on a PE use four parallel worker threads. Our system can use the four remaining hardware threads on each PE in order to keep disturbance of the actual computation at a minimum. Our software and experimental data are available at <https://github.com/domschrei/mallob>.

5.1 Uniform Jobs

In a first set of experiments, we analyze the base performance of our system by introducing a stream of jobs in such a way that exactly n_{par} jobs are in the system at any time. We limit each job j to a CPU time budget B inversely proportional to n_{par} . Each job corresponds to a difficult SAT formula which cannot be solved within the given budget. As such, we emulate jobs of fixed size.

We chose m and the values of n_{par} in such a way that $m/n_{\text{par}} \in \mathbb{N}$ for all runs. We compare our runs against a hypothetical rigid scheduler which functions as follows: Exactly m/n_{par} PEs are allotted for each job, starting with the first n_{par} jobs at $t = 0$. At periodic points in time, all jobs finish and each set of PEs instantly receives the next job. This leads to perfect utilization and maximizes throughput. We neglect any kind of overhead for this scheduler.

For a modest number of parallel jobs n_{par} in the system ($n_{\text{par}} \leq 192$), our scheduler reaches 99% of the optimal rigid scheduler's throughput (Table 1). This efficiency decreases to 97.6% for the largest n_{par} where $v_j = 2$ for each job. As the CPU time of each job is calculated in terms of its assigned volume and as the allocation of workers takes some time, each job uses slightly less CPU time than advertised: Dividing the time for which each job's workers have been active by its advertised CPU time, we obtained a work efficiency of $\eta \geq 99\%$. Lastly, we measured the CPU utilization of all worker threads as reported by the operating system, which averages at 98% or more. In terms of overall work efficiency $\eta \times u$, we observed an optimum of 98% at $n_{\text{par}} = 192$, a point where neither n_{par} nor the size of individual job trees is close to m .

Table 1. Scheduling uniform jobs on 1536 PEs (6144 cores) compared to a hypothetical optimal rigid scheduler. From left to right: Max. number n_{par} of parallel jobs; max. measured throughput θ , optimal throughput θ_{opt} (in jobs per second), throughput efficiency $\theta/\theta_{\text{opt}}$; work efficiency η ; mean measured CPU utilization u of worker threads.

n_{par}	θ	θ_{opt}	$\frac{\theta}{\theta_{\text{opt}}}$	η	u
3	0.159	0.16	0.991	0.990	0.981
6	0.318	0.32	0.994	0.990	0.983
12	0.636	0.64	0.993	0.991	0.984
24	1.271	1.28	0.993	0.992	0.985
48	2.543	2.56	0.993	0.993	0.985
96	5.071	5.12	0.990	0.993	0.986
192	10.141	10.24	0.990	0.995	0.985
384	20.114	20.48	0.982	0.995	0.983
768	39.972	40.96	0.976	0.992	0.980

5.2 Impact of Priorities

In the following we evaluate the impact of job priorities. We use 32 nodes (1536 cores, 384 PEs) and introduce nine streams of jobs, each stream with a different job priority $p \in [0.01, 1]$ (see Fig. 4 right) and with a wallclock limit of 300 s per job. As such, the system processes nine jobs with nine different priorities at a time. Each stream is a permutation of 80 diverse SAT instances [27].

As expected, we observed a proportional relationship between priority and assigned volume, with small variations due to rounding (Fig. 4). By contrast, response times appear to decrease exponentially towards a certain lower bound, which is in line with the NP-hardness of SAT and the diminishing returns of parallel SAT solving [27]. The modest margin by which average response times decrease is due to the difficulty of the chosen SAT benchmarks, many of which cannot be solved within the imposed time limit at either scale.




Fig. 4. Impact of job priority on mean assigned volume (left axis, blue triangles) and response time (right axis, orange squares). The table shows the used priorities p_j with the corresponding mean assigned volume \tilde{v}_j and mean response times in seconds. (Color figure online)

5.3 Realistic Job Arrivals

In the next set of experiments, we analyze the properties of our system in a more realistic scenario. Four PEs introduce batches of jobs with poisson-distributed arrivals (inter-arrival time of $1/\lambda \in \{2.5\text{ s}, 5\text{ s}, 10\text{ s}\}$) and between one and eight jobs per batch. As such, we simulate users which arrive at independent times and submit a number of jobs at once. We also sample a priority $p_j \in [0.01, 1]$, a maximum demand $d_j \in 1, \dots, 1536$, and a wallclock limit $b_j \in [1, 600]\text{ s}$ for each job. We ran this experiment with our current request matching (Sect. 4.1) and with each request message performing up to h random hops (as in [27]) until our request matching is employed, for varying values of h . In addition, we ran the experiment with three different suspended worker reuse strategies: No deliberate reuse at all, the basic approach from [27], and our current approach.

Figure 5(left) shows the number of active jobs in the system over time for our default configuration (our reuse strategy and immediate matching of requests). For all tested interarrival times, considerable changes in the system load can be observed during a job's average life time which justify the employment of a malleable scheduling strategy. Figure 5(right) illustrates for $1/\lambda = 5\text{ s}$ that system utilization is at around 99.8% on average and almost always above 99.5%. We also measured the ratio of time for which each PE has been idle: The median PE was busy 99.08% of all time for the least frequent job arrivals ($1/\lambda = 10\text{ s}$), 99.77% for $1/\lambda = 5\text{ s}$, and 99.85% for $1/\lambda = 2.5\text{ s}$. Also note that $\sum_j d_j < m$ for the first seconds of each run, hence not all PEs can be utilized immediately.

In the following, we focus on the experiment with $1/\lambda = 5\text{ s}$. The latency of our volume calculation, i.e., the latency until a PE received an updated volume for an updated job, reached a median of 1 ms and a maximum of 34 ms for our default configuration. For the scheduling of an arriving job, Fig. 6(left) shows that the lowest latencies were achieved by our request matching ($h = 0$). For increasing values of h , the variance of latencies increases and high latencies ($\geq 50\text{ ms}$) become more and more likely. Note that jobs normally enter a fully utilized system, and have $d_j = 1$. Therefore, the triggered balancing calculation may render only a single PE idle, which heavily disfavors performing a random walk. Regarding the latency of expanding a job tree by another layer,




Fig. 5. Left: Number of active jobs for interarrival times $1/\lambda$ of 2.5 s (top), 5 s (middle), and 10 s (bottom). Right: System utilization (i.e., ratio of busy PEs) for $1/\lambda = 5\text{ s}$ at a sliding average of window size 1 s, 15 s, and 60 s respectively.




Fig. 6. Distribution over measured latency for the initial scheduling of a job (left) and the expansion of a job tree by another worker (right), for inter arrival rate $1/\lambda = 5$, for a varying number h of random hops until a request message is routed along R .

Fig. 6(right) indicates that requests performing random walks have a high chance to succeed quickly but can otherwise result in high latencies (>10 ms).

To compare strategies for reusing suspended workers, we divided the number of created workers for a job j by its maximum assigned volume v_j . This *Worker Creation Ratio* (WCR) is ideally 1 and becomes larger the more often a worker is suspended and then re-created at a different PE. We computed the WCR for each job and in total: As Table 2 shows, our approach reduces a WCR of 2.14 down to 1.8 (-15.9%). Context switches (i.e., how many times a PE changed its affiliation) and average response times are improved marginally compared to the naive approach. Last but not least, we counted on how many distinct PEs each w_j^i has been created: Our strategy initializes 89% of all workers only once, and 94% of workers have been created at most five times. We conclude that most jobs only feature a small number of workers which are rescheduled frequently.

Table 2. Comparison of worker reuse strategies in terms of worker creation ratio (WCR, per job – median, maximum – and in total), context switches (CS, median per PE and mean), the number of processed jobs within 1 h (Pr.), their mean response time (RT), and the fraction of workers created on at most $\{1, 2, 5, 10, 25\}$ distinct PEs.

	WCR			CS		Pr.	RT	Pr [WC \leq .]				
	med.	max.	total	med.	mean			1	2	5	10	25
None	1.43	33.0	2.14	136	138.2	5923	153.40	0.87	0.90	0.94	0.97	0.992
Basic	1.40	31.5	2.07	134	135.3	5921	153.89	0.87	0.90	0.94	0.97	0.993
Ours	1.25	24.5	1.80	130	131.8	5939	152.33	0.89	0.91	0.94	0.97	0.993

6 Conclusion

We have presented a decentralized and highly scalable approach to online job scheduling of malleable NP-hard jobs with unknown processing times. We split our problem into two subproblems, namely the computation of fair job volumes

and the assignment of jobs to PEs, and proposed scalable distributed algorithms with $\mathcal{O}(\log m)$ span for both of them. We presented a practical implementation and experimentally showed that it schedules incoming jobs within tens of milliseconds, distributes resources proportional to each job's priority, and leads to near-optimal utilization of resources.

For future work, we intend to add engines for applications beyond SAT into our system. Furthermore, we want to generalize our approach to heterogeneous computing environments and add fault tolerance to our distributed algorithms.

Acknowledgments and Data Availability. The datasets and code generated and/or analyzed during this study are available in the Figshare repository: <https://doi.org/10.6084/m9.figshare.20000642> [23]. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 882500). The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer SuperMUC-NG at Leibniz Supercomputing Centre (www.lrz.de). The authors wish to thank Tim Niklas Uhl as well as the anonymous reviewers for their helpful feedback.



References

1. Ajtai, M., Komlós, J., Szemerédi, E.: Sorting in $\log n$ parallel steps. *Combinatorica* **3**(1), 1–19 (1983). <https://doi.org/10.1109/tc.1985.5009385>
2. Alquraan, A., Takruri, H., Alfatafta, M., Al-Kiswany, S.: An analysis of network-partitioning failures in cloud systems. In: *Symposium on Operating Systems Design and Implementation*, pp. 51–68 (2018)
3. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. In: *International Joint Conference on Artificial Intelligence*, pp. 399–404 (2009)
4. Axtmann, M., Sanders, P.: Robust massively parallel sorting. In: *Meeting on Algorithm Engineering and Experiments (ALENEX)*, pp. 83–97 (2017). <https://doi.org/10.1137/1.9781611974768.7>
5. Blazewicz, J., Kovalyov, M.Y., Machowiak, M., Trystram, D., Weglarz, J.: Preemptible malleable task scheduling problem. *IEEE Trans. Comput.* **55**(4), 486–490 (2006). <https://doi.org/10.1109/tc.2006.58>
6. Buisson, J., Sonmez, O., Mohamed, H., Lammers, W., Epema, D.: Scheduling malleable applications in multicluster systems. In: *International Conference on Cluster Computing*, pp. 372–381. IEEE (2007). <https://doi.org/10.1109/clustr.2007.4629252>
7. Castro, M., Liskov, B.: Practical byzantine fault tolerance. In: *Symposium on Operating Systems Design and Implementation*. pp. 173–186 (1999)
8. Cook, S.A.: The complexity of theorem-proving procedures. In: *ACM symposium on Theory of Computing*, pp. 151–158 (1971). <https://doi.org/10.7551/mitpress/12274.003.0036>
9. Desell, T., El Maghraoui, K., Varela, C.A.: Malleable applications for scalable high performance computing. *Clust. Comput.* **10**(3), 323–337 (2007). <https://doi.org/10.1007/s10586-007-0032-9>

10. Feitelson, D.G.: Job scheduling in multiprogrammed parallel systems (1997)
11. Froleys, N., Heule, M., Iser, M., Jarvisalo, M., Suda, M.: SAT competition 2020. *Artif. Intell.* **301**, 103572 (2021). <https://doi.org/10.1016/j.artint.2021.103572>
12. Gropp, W., Lusk, E., Skjellum, A.: Using MPI: Portable Parallel Programming with the Message-Passing Interface, vol. 1. MIT Press, London (1999). <https://doi.org/10.7551/mitpress/7056.001.0001>
13. Gupta, A., Acun, B., Sarood, O., Kalé, L.V.: Towards realizing the potential of malleable jobs. In: International Conference on High Performance Computing (HiPC), pp. 1–10. IEEE (2014). <https://doi.org/10.1109/hipc.2014.7116905>
14. Hamadi, Y., Jabbour, S., Sais, L.: ManySAT: a parallel SAT solver. *J. Satisf. Boolean Model. Comput.* **6**(4), 245–262 (2010). <https://doi.org/10.3233/sat190070>
15. Heisinger, M., Fleury, M., Biere, A.: Distributed cube and conquer with Paracooba. In: Pulina, L., Seidl, M. (eds.) SAT 2020. LNCS, vol. 12178, pp. 114–122. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-51825-7_9
16. Huang, C., Lawlor, O., Kalé, L.V.: Adaptive MPI. In: Rauchwerger, L. (ed.) LCPC 2003. LNCS, vol. 2958, pp. 306–322. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24644-2_20
17. Hungershofer, J.: On the combined scheduling of malleable and rigid jobs. In: Symposium on Computer Architecture and HPC, pp. 206–213. IEEE (2004). <https://doi.org/10.1109/sbac-pad.2004.27>
18. Kleine Büning, M., Balyo, T., Sinz, C.: Using DimSpec for bounded and unbounded software model checking. In: Ait-Ameur, Y., Qin, S. (eds.) ICFEM 2019. LNCS, vol. 11852, pp. 19–35. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32409-4_2
19. Marques-Silva, J., Lynce, I., Malik, S.: Conflict-driven clause learning SAT solvers. In: Handbook of Satisfiability, pp. 131–153. IOS Press (2009). <https://doi.org/10.3233/faia200987>
20. Massacci, F., Marraro, L.: Logical cryptanalysis as a SAT problem. *J. Autom. Reason.* **24**(1), 165–203 (2000). <https://doi.org/10.1023/A:1006326723002>
21. Ozdemir, A., Wu, H., Barrett, C.: SAT solving in the serverless cloud. In: Formal Methods in Computer Aided Design (FMCAD), pp. 241–245. IEEE (2021). https://doi.org/10.34727/2021/isbn.978-3-85448-046-4_33
22. Sanders, P., Mehlhorn, K., Dietzfelbinger, M., Dementiev, R.: Sequential and Parallel Algorithms and Data Structures: The Basic Toolbox. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-25209-0>
23. Sanders, P., Schreiber, D.: Artifact and instructions to generate experimental results for the Euro-Par 2022 paper: “Decentralized Online Scheduling of Malleable NP-hard Jobs”. <https://doi.org/10.6084/m9.figshare.20000642>
24. Sanders, P., Speck, J.: Efficient parallel scheduling of malleable tasks. In: International Parallel and Distributed Processing Symposium, pp. 1156–1166. IEEE (2011). <https://doi.org/10.1109/ipdps.2011.110>
25. Sanders, P., Speck, J.: Energy efficient frequency scaling and scheduling for malleable tasks. In: Kaklamanis, C., Papatheodorou, T., Spirakis, P.G. (eds.) Euro-Par 2012. LNCS, vol. 7484, pp. 167–178. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32820-6_18
26. Schreiber, D.: Lilotane: a lifted SAT-based approach to hierarchical planning. *J. Artif. Intell. Res.* **70**, 1117–1181 (2021). <https://doi.org/10.1613/jair.1.12520>
27. Schreiber, D., Sanders, P.: Scalable SAT solving in the cloud. In: Li, C.-M., Manyà, F. (eds.) SAT 2021. LNCS, vol. 12831, pp. 518–534. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-80223-3_35

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

