



Scheduling computations with provably low synchronization overheads

Guilherme Rito¹ · Hervé Paulino²

Accepted: 10 August 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We present a Work Stealing scheduling algorithm that provably avoids most synchronization overheads by keeping processors' dequeues entirely private by default and only exposing work when requested by thieves. This is the first paper that obtains bounds on the synchronization overheads that are (essentially) independent of the total amount of work, thus corresponding to a great improvement, in both algorithm design and theory, over state-of-the-art Work Stealing algorithms. Consider any computation with work T_1 and critical-path length T_∞ executed by P processors using our scheduler. Our analysis shows that the expected execution time is $O\left(\frac{T_1}{P} + T_\infty\right)$, and the expected synchronization overheads incurred during the execution are at most $O((C_{CAS} + C_{MFence})PT_\infty)$, where C_{CAS} and C_{MFence} , respectively, denote the maximum cost of executing a Compare-And-Swap instruction and a Memory Fence instruction.

Keywords Work stealing · Synchronization overheads · Upper bounds · Scheduling

1 Introduction

In Work Stealing, each worker (usually referred to as *processor*) owns a double-ended queue (deque) of threads ready to execute. This deque is locally manipulated as a stack, similar to a sequential execution: Processors push and pop threads from the bottom side of their deque when, respectively, a new thread is spawned and the execution of the current thread concludes. Additionally, whenever a pop operation finds the local deque empty, the processor becomes a *thief* and starts targeting other processors—called its *victims*—uniformly at random, with the purpose of stealing a thread from the top of their dequeues.

Work Stealing is a provably efficient algorithm for scheduling multithreaded computations (Blumofe & Leiserson, 1999). However, as shown by Attiya et al., (2011), due

to the concurrent nature of processors' dequeues, the use of appropriate synchronization mechanisms in these data structures is required for correctness. Consequently, even when processors are operating locally on their dequeues, they incur expensive synchronization overheads that, in most cases, are unnecessary.

The first provably efficient Work Stealing algorithm, proposed in Blumofe & Leiserson, (1999), assumed that all steal attempts targeting each deque were serialized, and only ensured the success of at most one such attempt per time step. The idea was materialized in Cilk (Blumofe et al., 1996) via a blocking synchronization protocol named *THE*. Despite being extremely efficient, Frigo et al. found that the overheads introduced by the *THE* protocol easily account for more than half of Cilk's total execution time (Frigo et al., 1998). Subsequent work mitigated part of these overheads by replacing the *THE* protocol with a non-blocking one that resorts to Compare-And-Swap (CAS) and Memory Fence (MFence) instructions (Arora et al., 2001; Blumofe & Papadopoulos, 1998), which are atomic instructions that are used to achieve synchronization between processors, in this case being used to synchronize concurrent deque accesses by processors. Later, Morrison et al. tuned Cilk by removing a single MFence instruction (one that was executed whenever a processor tried to take work from its deque) and found that this single MFence could account for as much as 25% of

✉ Hervé Paulino
herve.paulino@fct.unl.pt

Guilherme Rito
guilherme.teixeira@inf.ethz.ch

¹ Department of Computer Science, ETH-Zurich, Zurich, Switzerland

² NOVA Laboratory for Computer Science and Informatics (NOVA LINCS), NOVA School of Science and Technology, NOVA University Lisbon, Caparica, Portugal

the total execution time (Morrison et al., 2014). Unfortunately, as proved in Attiya et al., (2011), it is impossible to eliminate all synchronization (e.g., the MFence instruction mentioned above) from the implementation of any concurrent data structure that could possibly be used as a work-queue by a Work Stealing algorithm, while maintaining correctness. Indirectly, this result implies the impossibility of eliminating all synchronization from Work Stealing algorithms that use any fully concurrent data structure as processor's work-queues.

Various proposals have been made with the goal of eliminating synchronization for local deque accesses, by making deques partly or even entirely private (Acar et al., 2013; Dinan et al., 2009; Hiraishi et al., 2009; Lifflander et al., 2012; Morrison et al., 2014; Tzannes et al., 2011; van Dijk and van de Pol, 2014). The elimination of synchronization for local deque accesses, however, raises a new problem. Since synchronization is required to guarantee correctness, when a processor p spawns a thread Γ and pushes Γ (locally) to its work-queue, Γ cannot safely be stolen from p by other processors, at least until p issues some synchronization operation. So, when should a busy processor use synchronization to permit load balancing? The subtlety of this question is evidenced by the inexistence of any algorithm that provably avoids most synchronization overheads while maintaining provably good performance. On the one hand, if a processor exposes work too eagerly, then it still incurs unnecessary synchronization overheads (Acar et al., 2013; Dinan et al., 2008, 2009; Lifflander et al., 2012; Tzannes et al., 2011; van Dijk & van de Pol, 2014). On the other hand, if a processor barely exposes any work, then load balancing opportunities become limited, thus potentially dropping the asymptotically optimal runtime guarantees of Work Stealing (Hiraishi et al., 2009; Morrison et al., 2014; van Dijk & van de Pol, 2014). To address this problem optimally, our algorithm follows a lazy approach: (1) A processor p only uses synchronization to expose work when a thief directly asks p for work, and (2) p only exposes a single unit of work (i.e., a single thread) for each time it is asked to expose work.

1.1 Contributions

In this paper, we present Low-Cost Work Stealing, a variant of the Work Stealing algorithm that uses split deques to provably avoid most synchronization overheads, while maintaining an asymptotically optimal expected runtime. The theoretical significance of our contributions is highlighted, for instance, by the tight bounds we obtain on the synchronization overheads incurred by our algorithm. Our bounds are essentially independent from the computation's total amount of work, contrasting with previous work. From an algorithm design perspective, Low-Cost Work Stealing greatly improves over prior Work Stealing schedulers as it shows how to optimally

use synchronization to permit provably efficient load balancing. Four of the distinctive features of our algorithm are:

1. Busy processors only expose work to be stolen after being targeted by one or more steal attempts. This allows processors to work locally on their work-queue without requiring any synchronization, imposing it only when load balancing may be needed.
2. Work exposure requests are attended in constant time. This is crucial to keep the algorithm's execution time bounds a constant factor away from optimal. The requirement may be achieved by periodically checking for requests or by implementing an asynchronous notification mechanism. For the sake of simplicity, we only focus on the former.
3. Processors only expose one thread of their local work at a time, contrasting with prior approaches. Consequently, synchronization for local operation can be mostly avoided when little load balancing is needed.
4. All interactions between processors are completely asynchronous, making our algorithm viable for multiprogrammed environments.

As we will see, our analysis shows that for a P -processor execution of a computation with total work T_1 and critical-path length (i.e., span) T_∞ , the expected execution time of Low-Cost Work Stealing is at most $O\left(\frac{T_1}{P} + T_\infty\right)$, and the expected synchronization overheads incurred by the algorithm are at most $O((C_{CAS} + C_{MFence})PT_\infty)$, where C_{CAS} and C_{MFence} , respectively, denote the synchronization costs incurred by the execution of a CAS and MFence instructions. These bounds are tight and imply that for several classes of computations our algorithm reduces the use of synchronization by an almost exponential factor when compared with prior provably efficient Work Stealing algorithms.

2 Preliminaries

Like in much previous work (Acar et al., 2002, 2013; Agrawal et al., 2007, 2008; Arora et al., 1998, 2001; Blumofe & Leiserson, 1999; Muller & Acar, 2016; Tchiboukdjian et al., 2010), we model a computation as a *dag* (i.e., a direct acyclic graph) $G = (V, E)$, where each node $v \in V$ corresponds to an instruction, and each edge $(\mu_1, \mu_2) \in E$ denotes an ordering between two instructions (meaning μ_2 can only be executed after μ_1). Nodes with in-degree of 0 are referred to as *roots*, while nodes with out-degree of 0 are called *sinks*. Equivalently to Arora et al., (2001), we make two assumptions related to the structure of computations. Let G denote a computation's dag:




Fig. 1 The Split Deque built from an array of nodes (named **entries**), and featuring a state composed of 3 variables: **age** comprising fields *top*, that points to the split deque's topmost node, and *tag*, to ensure correctness (see Sect. 3.3); **official bottom** points to node below the

bottommost one of the split deque's public part, and, lastly, **private bottom** points to the empty slot below the split deque's bottommost node

1. there exists only one root and one sink in G ;
2. the out-degree of any node within G is at most two (meaning that each instruction can spawn at most one thread).

The total number of nodes within a dag is expressed by T_1 and the length of a longest directed path (i.e., the critical-path length) by T_∞ . A node is *ready* if all its ancestors have been executed, implying that all the ordering constraints of E are satisfied. When a node becomes ready, we say that it was *enabled*; to ensure correctness only ready nodes can be executed. The assignment of a node μ to a processor p means that μ will be the next node p executes. Finally, a computation's execution can be partitioned into discrete time steps, such that at each step, every processor executes an instruction.

3 Low-Cost Work Stealing

In Low-Cost Work Stealing, each processor owns a Lock-Free split deque, instead of a typical concurrent deque. A split deque (illustrated in Fig. 1) is simply a deque that is split into two parts: a private part and a public part. The public part lies in the top of the split deque whereas the private corresponds to the rest of the split deque. To avoid synchronization for local operations, only the owner of a split deque is allowed to access its private part. Furthermore, by default busy processors operate on the private part of their split deque, pushing and popping ready nodes as necessary. In fact, a busy processor only attempts to fetch work from the public part of its split deque if the private part is empty. In such situation, if the processor's attempt succeeds (i.e., if the public part of the processor's split deque is not empty), the obtained node becomes the processor's new assigned node. However, if the public part of the processor's split deque is also empty, the processor becomes a thief and begins a stealing phase. During stealing phases, thieves target victims

uniformly at random and attempt to steal work from the top of their split deques. To keep the private part of split deques entirely private, steal attempts are only allowed to access the public part. Thus, when a thief attempts to steal work from a victim's split deque whose public part is empty (illustrated in Fig. 1a), the steal attempt simply fails and the thief does not obtain work. In that case, the thief then updates a victim's flag (referred to as the *targeted* flag) to (asynchronously) notify the victim that the public part of its split deque is empty (more on this in Sect. 3.2). When the owner of the split deque realizes it was notified (by checking the value of its *targeted* flag), it tries to transfer a node from the private part of its split deque to the public part. If the private part is not empty, then a node is transferred, in which case we say that the transferred node became *stealable* (illustrated in Fig. 1b).

3.1 The lock-free Split Deque

We now present the specification of a split deque object, along with its associated relaxed semantics. Being the behavior of split deques similar to the behavior of concurrent deques, the split deque's relaxed semantics are comparable to the relaxed deque semantics presented in Arora et al., (2001). A split deque object meeting the relaxed semantics supports five methods:

- PUSH—Pushes a node into the bottom of the split deque's private part.
- POP—Removes and returns a node from the bottom of the split deque's private part, if that part is not empty. Otherwise, returns the special value RACE.
- UPDATEBOTTOM—Transfers the topmost node from the private part of the split deque into the bottom of the public part, and does not return a value. The invocation has no effect if the private part of the split deque is empty.

POPBOTTOM—Removes and returns the bottom-most node from the public part of the split deque. If this part of the split deque is empty, the invocation has no effect and EMPTY is returned.

POPTOP—Attempts to remove and return the topmost node from the public part of the split deque. If the public part is empty, the invocation has no effect and the value EMPTY is returned. If the invocation aborts, it has no effect and the value ABORT is returned.

A split deque implementation is constant-time *iff* any invocation to each of these methods takes at most a constant number of steps to return. Say that a set of invocations to a split deque's methods meets the relaxed semantics *iff* there is a set of *linearization times* for the corresponding non-aborting invocations such that:

1. Every non-aborting invocation's linearization time lies within the beginning and completion times of the respective invocation;
2. No linearization times associated with distinct non-aborting invocations coincide;
3. The return values for the non-aborting invocations are consistent with a serial execution of the methods in the order given by the linearization times of the corresponding non-aborting invocations; and
4. For each aborted POPTOP invocation x to a split deque d , there exists another invocation removing the topmost item from d whose linearization time falls between the beginning and completion times of invocation x .

3.2 The Low-Cost Work Stealing algorithm

Algorithm 1 depicts the specification of the Low-Cost Work Stealing algorithm. Each processor owns a split deque that it uses to store its attached nodes and, additionally, owns a *targeted* flag that stores a Boolean value. This flag is used to implement an asynchronous notification mechanism that allows thieves to request their victims to expose work, allowing it to be stolen. Even though, in practice, the notification mechanism of our algorithm can be implemented using signals, to perform a correct analysis of the algorithm's synchronization overheads all the possible sources of such overheads must be explicit, for which reason we chose to embed a simple notification mechanism into the algorithm's specification. Although the *targeted* flag of each processor can be simultaneously accessed by multiple processors, to ensure the algorithm's correctness it suffices to guarantee that busy processors read an up-to-date value of their *targeted* flag, an operation not requiring any synchronization (see Sewell et al., 2010).

Algorithm 1 Low-Cost Work Stealing algorithm.

```

1: procedure SCHEDULER
2:   while computation not terminated do
3:     if self.targeted then
4:       self.spdeque.UPDATEBOTTOM()
5:       self.targeted  $\leftarrow$  FALSE
6:     end if
7:     if VALIDNODE(assigned) then
8:       enabled  $\leftarrow$  EXECUTE(assigned)
9:       if LENGTH(enabled) > 0 then
10:        assigned  $\leftarrow$  enabled[0]
11:        if LENGTH(enabled) = 2 then
12:          self.spdeque.PUSH(enabled[1])
13:        end if
14:       else
15:        assigned  $\leftarrow$  self.spdeque.POP()
16:        if assigned = RACE then
17:          assigned  $\leftarrow$  self.spdeque.POPBOTTOM()
18:        end if
19:       end if
20:     else
21:       self.WORKMIGRATION()
22:     end if
23:   end while
24: end procedure

25: procedure WORKMIGRATION
26:   victim  $\leftarrow$  UNIFORMLYRANDOMPROCESSOR()
27:   assigned  $\leftarrow$  victim.spdeque.POPTOP()
28:   if assigned = EMPTY then
29:     victim.targeted  $\leftarrow$  TRUE
30:   end if
31: end procedure

32: function VALIDNODE(node)
33:   return node  $\neq$  EMPTY and node  $\neq$  ABORT and node  $\neq$ 
      NONE
34: end function

```

Before a computation's execution begins, every processor sets its *assigned* node to NONE and its *targeted* flag to FALSE. To start the execution, one of the processors gets the root node assigned.

As we will see, the behavior of Low-Cost Work Stealing is similar to the original Work Stealing algorithm. Consider some processor p working on a computation scheduled by Low-Cost Work Stealing, and some iteration of the scheduling loop that p executes (corresponding to lines 2 to 23 of Algorithm 1). First, p reads the value of its *targeted* flag to check if it has been notified by some thief. If p 's *targeted* flag is set to TRUE (i.e., if p was notified), the processor tries to make a node stealable, by invoking UPDATEBOTTOM to its split deque. After that, and regardless of that invocation's outcome, p resets its *targeted* flag back to FALSE. The subsequent behavior of p depends on whether it has an assigned node.

- If p has an assigned node, p executes the node. From this execution, either zero, one or two nodes can be enabled.

Zero nodes enabled: The processor tries to fetch the bottommost node stored in its split deque. To that end, p first tries to fetch a node from the bottom of its split deque’s private part (line 15). If p finds that part empty, it then tries to fetch a node from the public part (line 17). If this part is also empty, p becomes a thief and starts a work stealing phase. On the other hand, if p successfully fetched a node from any of the parts of its split deque, then the returned node becomes p ’s new assigned node.

One node enabled: The enabled node becomes p ’s new assigned node (line 10).

Two nodes enabled: One of the enabled nodes becomes p ’s new assigned node, while the other is pushed into the bottom of the private part of p ’s split deque (line 12).

- If p does not have an assigned node, it is searching for work. In this situation, the processor first targets, uniformly at random, a victim processor and then attempts to steal work from the public part of the victim’s split deque (lines 26 and 27). If the attempt is successful, the stolen node becomes p ’s new assigned node. If the attempt aborts, p simply gives up on the steal attempt. For last, if p finds the public part of the victim’s split deque empty it sets the victim’s *targeted* flag to TRUE (line 29), notifying the victim that it found the public part of the victim’s split deque empty.

3.3 A Split Deque implementation

Algorithm 2 depicts a possible implementation of the lock-free split deque, based on the deque’s implementation given in Arora et al., (2001). As illustrated in Fig. 1, each split deque object has four instance variables:

entries—an array of ready nodes.

privateBottom—the index below the bottommost node of the split deque.

officialBottom—the index below the bottommost node of the split deque’s public part.

age—composed of two fields: *top*, the index of the top node, and *tag*, which is only used to ensure correction (avoiding the ABA problem).

Algorithm 2 The Split Deque implementation.

```

privateBottom ← 0                                ▷ private field
entries ← {}                                       ▷ private read-write, public read-only
officialBottom ← 0                                 ▷ private read-write, public read-only
age ← {0, 0}                                       ▷ public field

1: procedure PUSH(node)
2:  $pBot \leftarrow self.privateBottom$ 
3:  $self.entries[pBot] \leftarrow node$ 
4:  $self.privateBottom \leftarrow pBot + 1$ 
5: end procedure

6: procedure POP
7:  $pBot \leftarrow self.privateBottom$ 
8: if  $pBot = self.officialBottom$  then return RACE
9: end if
10:  $pBot \leftarrow pBot - 1$ 
11:  $node \leftarrow self.entries[pBot]$ 
12:  $self.privateBottom \leftarrow pBot$ 
13: return node
14: end procedure

15: procedure POPTOP
16:  $oldAge \leftarrow self.age$ 
17:  $oldBottom \leftarrow self.officialBottom$ 
18: if  $oldBottom \leq oldAge.top$  then return EMPTY
19: end if
20:  $node \leftarrow self.entries[oldAge.top]$ 
21:  $newAge \leftarrow oldAge$ 
22:  $newAge.top \leftarrow newAge.top + 1$ 
23: if CAS(age, oldAge, newAge) = SUCCESS then
24:   return node
25: end if
26: return ABORT
27: end procedure

28: procedure UPDATEBOTTOM
29:  $pBot \leftarrow self.privateBottom$ 
30:  $oBot \leftarrow self.officialBottom$ 
31: if  $pBot > oBot$  then  $oBot \leftarrow oBot + 1$ 
32: end if
33:  $self.officialBottom \leftarrow oBot$ 
34: end procedure

35: procedure POPBOTTOM
36:  $oBot \leftarrow self.officialBottom$ 
37: if  $oBot = 0$  then return EMPTY
38: end if
39:  $oBot \leftarrow oBot - 1$ 
40:  $self.officialBottom \leftarrow oBot$ 
41:  $node \leftarrow self.entries[oBot]$ 
42:  $oldAge \leftarrow age$ 
43: if  $oBot > oldAge.top$  then return node
44: end if
45:  $self.officialBottom \leftarrow 0$ 
46:  $self.privateBottom \leftarrow 0$ 
47:  $newAge.top \leftarrow 0$ 
48:  $newAge.tag \leftarrow oldAge.tag + 1$ 
49: if  $oBot = oldAge.top$  then
50:   if CAS(age, oldAge, newAge) = SUCCESS then
51:     return node
52:   end if
53: end if
54:  $self.age \leftarrow newAge$ 
55: return EMPTY
56: end procedure

```


We say that a set of invocations is *good* if and only if the methods PUSH, POP, UPDATEBOTTOM and POPBOTTOM are never invoked concurrently. For Low-Cost Work Stealing, as only the owner of each split deque can invoke these methods, it is easy to deduce that all sets of invocations issued by the algorithm are good. Furthermore, we claim that the implementation depicted in Algorithm 2 is constant-time and meets the relaxed semantics (defined in Sect. 3.1) on any good set of invocations. However, even though all methods are composed by a small number of instructions and none includes a loop, proving this claim is not a straightforward task because all possible execution interleavings have to be considered. Moreover, as the main focus of this study is not related to programs' verification, the proof of this claim falls out of the scope of this paper. Yet, we remark that the proposed implementation is a simple extension of the deque implementation presented in Arora et al., (2001), which has been proven in Blumofe and Leiserson (1999) to be a correct implementation, meeting the relaxed deque semantics on any set of invocations made by the Work Stealing algorithm. For this reason, throughout this paper we assume that for any set of invocations issued by the Low-Cost Work Stealing algorithm, the relaxed semantics is always satisfied.

Lemma 1 *No invocation to PUSH requires an MFence instruction.*

Proof Since the PUSH method operates only once over a single publicly accessible field (*entries*) of the split deque's state, no MFence instructions are required. \square

Lemma 2 *No invocation to POP requires an MFence instruction.*

Proof Any invocation to the POP method only reads from two publicly accessible fields of the split deque's state, namely *officialBottom* (line 8) and *entries* (line 12). However, due to a data dependency, no re-ordering between these read operations may occur, and so, no MFence instructions are required. \square

The dag of a computation is dynamically unfolded during its execution. If the execution of a node u enables another node u' , then (u, u') is an *enabling edge* and refer to node u as the *designated parent* of u' . Refer to the tree formed by the enabling edges of a particular execution of a dag by *enabling tree*, and denote the depth of a node u within this tree by $d(u)$. Define the weight of u as $w(u) = T_\infty - d(u)$. Similar to Arora et al., (2001), our analysis is made in an *a posteriori* fashion, allowing us to refer to the enabling tree generated by a computation's execution.

The next lemma states the standard structural property of deques.

Lemma 3 (Structural Lemma for split deques) *Let v_1, \dots, v_k denote the nodes stored in some processor p 's split deque,*

ordered from the bottom of the split deque to the top, at some point in the linearized execution of Low-Cost Work Stealing. Moreover, let v_0 denote p 's assigned node (if any), and for $i = 0, \dots, k$ let u_i denote the designated parent of v_i in the enabling tree. Then, for $i = 1, \dots, k$, u_i is an ancestor of u_{i-1} in the enabling tree, and although v_0 and v_1 may have the same designated parent (i.e., $u_0 = u_1$), for $i = 2, 3, \dots, k$, $u_{i-1} \neq u_i$ (i.e., the ancestor relationship is proper).

Corollary 1 *Let v_1, \dots, v_k denote the nodes stored in some processor p 's split deque, ordered from the bottom of the split deque to the top, at some moment during the execution of Low-Cost Work Stealing. Moreover, let v_0 denote p 's assigned node (if any). Then, we have $w(v_0) \leq w(v_1) < \dots < w(v_{k-1}) < w(v_k)$.*

4 Analysis

In this section, we obtain bounds on the expected execution time of computations using Low-Cost Work Stealing and on the expected synchronization overheads incurred by the scheduler. The analysis we make follows the same overall idea as the one given in Arora et al., (2001). For the sake of readability, most of the proofs are deferred to "Appendix A." We now introduce a few more definitions needed for the analysis.

Define a *scheduling iteration* as a sequence of instructions executed by a processor corresponding to a particular iteration of the scheduling loop (lines 2 to 23 of Algorithm 1). Thus, the full sequence of instructions executed by each processor during a computation's execution can be partitioned into scheduling iterations. As in Arora et al., (2001), we introduce the concept of a *milestone*: An instruction within the sequence executed by a processor is a milestone *iff* it corresponds to a node's execution (line 8) or to the return of a call to WORKMIGRATION (line 31). Taking into account the definition of a scheduling iteration, it is clear that any scheduling iteration of the algorithm includes a milestone. Refer to iterations whose milestone corresponds to a node's execution as *busy iterations*, and refer to the remainder as *idle iterations*. As one might note, if a processor has an assigned node at the beginning of an iteration's execution, the iteration is a busy one, and, otherwise, the iteration is an idle one. By observing the scheduling loop (lines 2 to 23 of Algorithm 1), and taking into account that the split deque's implementation is constant time, it is clear that any scheduling iteration is composed of a constant number of instructions. It then follows that any processor executes at most a constant number of instructions between two consecutive milestones. Throughout the analysis, let C denote a constant that is large enough to guarantee

that any sequence of instructions executed by a processor with length at least C includes a milestone.

We can now bound the execution time of a computation depending on the number of idle iterations that take place during that computation's execution. The proof of the following result is a trivial variant of Arora et al., (2001), Lemma 5 but considering the Low-Cost Work Stealing algorithm.

Lemma 4 *Consider any computation with work T_1 being executed by P processors, under Low-Cost Work Stealing. The execution time is $O\left(\frac{T_1}{P} + \frac{I}{P}\right)$, where I denotes the number of idle iterations executed by processors.*

As we will see, the following two results are key, as they show that the synchronization overheads incurred by Low-Cost Work Stealing (essentially) only depend on the number of idle iterations that take place during a computation's execution.

Lemma 5 *Consider a processor p executing a busy iteration such that p 's targeted flag is set to FALSE when p checks it at the beginning of the iteration. If the execution of p 's assigned node enables one or more nodes or if the private part of p 's split deque is not empty, then no MFence instruction is required during the execution of the iteration.*

Proof Consider the Low-Cost Work Stealing algorithm, depicted in Algorithm 1. The first action processor p takes for the execution of that iteration is checking the value of its *targeted* flag (line 3). Since, by the statement of this lemma, p 's *targeted* flag is set to FALSE at the moment when p checks the flag's value, p does not enter the *then* branch of the if statement. Moreover, as a consequence of the conditional statement of line 3, there is a control dependency that does not allow the instructions succeeding the conditional expression to be reordered with the evaluation of the condition, implying no MFence instruction is required until this point.

After that, p checks if it has an assigned node (line 7). Again, since the next action p takes depends on its currently assigned node, there is a control dependency from the instruction where p checks if it has a currently assigned node to both branches of the if statement. Thus, no instruction reordering between the evaluation of the condition and any of the instructions succeeding that evaluation can be made, implying no MFence instruction is required until here.

Because we assumed p was executing a busy iteration, by the definition of a busy iteration, p must have an assigned node. Hence, p executes its assigned node. Since the next action p takes (line 9) depends on the outcome of that node's execution, there is a control dependency between the execution of p 's assigned node and the execution of the sequence of instructions corresponding to each of the possible outcomes. Hence, no instruction reordering can be made, implying no MFence instruction is necessary until this point.

From that node's execution, three outcomes are possible:

0 nodes enabled In this case, p invokes the POP method on its own split deque (line 15). By Lemma 2, the invocation does not require the execution of a MFence instruction. Furthermore, the next instruction (line 16) has a data dependency on the value of p 's assigned node, for which reason it cannot be reordered with the invocation of the POP method and so no MFence instruction is required. Since we have assumed that the private part of p 's split deque was not empty, it is trivial to conclude that the POP invocation returns a node, which immediately becomes p 's new assigned node. Thus, after having a new node assigned p takes no further action during the iteration, meaning no MFence instruction was required for the execution of the iteration in this situation.

1 node enabled In this case the enabled node becomes p 's new assigned node (line 10). Next, p checks the number of nodes that were enabled. The assignment of one of the enabled nodes and the instruction where p checks the number of nodes enabled can be reordered. Fortunately, because *enabled* is a local variable (line 8) that is solely accessed by p , there is no harm for a parallel execution if the instructions are reordered and so no MFence instruction is required for this case as well. Because p enabled a single node it takes no further action during the iteration, implying the lemma holds in this situation as well.

2 nodes enabled Finally, for this case one of the enabled nodes becomes p 's new assigned node (line 10). Using the same reasoning as for the case where a single node was enabled, we conclude that no MFence instruction is required at least until the evaluation of the conditional statement of line 11. Because p enabled two nodes, p enters the *then* branch of the conditional statement and pushes the node it did not assign into the bottom of its split deque, by invoking the PUSH method (line 12). Since there is a control dependency between the execution of this instruction and the evaluation of the condition, no instruction reordering is allowed. Thus, no MFence instruction is required between these two instructions.

Finally, Lemma 1 states that an invocation to the PUSH method does not require a MFence instruction to be executed. Because after the invocation p takes no further action during the iteration, we deduce the lemma holds, concluding its proof. \square

The following lemma is a consequence of Lemma 5 and states that the number of CAS and MFence instructions executed during a computation's execution using Low-Cost Work Stealing only depends on the number of idle iterations and processors.

Lemma 6 Consider any computation being executed by the Low-Cost Work Stealing algorithm, using P processors. The number of Compare-And-Swap (CAS) and Memory Fence (MFence) instructions executed by processors during the computation's execution is at most $O(I + P)$, where I denotes the total number of idle iterations executed by processors.

Proof By observing Algorithms 1 and 2, it is easy to see that only invocations to POPBOTTOM or POPTOP methods can lead to the execution of CAS instructions. Furthermore, both these methods are invoked at most once per scheduling iteration, and, for both, at most one CAS instruction is executed per invocation. Since processors only invoke the POPTOP method when executing idle iterations, the number of CAS instructions caused by invocations to POPTOP is $O(I)$. On the other hand, processors only invoke the POPBOTTOM method during busy iterations where the private part of their split deque is empty and the execution of their currently assigned node does not enable any new nodes. Let p denote some processor executing one such iteration. From p 's invocation to the POPBOTTOM method two outcomes are possible:

A node is returned In this case the public part of p 's split deque was not empty implying p had previously transferred a node from the private part of its split deque to the public part. By observing Algorithm 1, it is easy to deduce that p only makes these node transfers if some thief had previously set p 's *targeted* flag to TRUE. Moreover, because after transferring the node p immediately sets its *targeted* flag back to FALSE, the number of times p makes such node transfers is at most the number of times it is targeted by a steal attempt. Taking into account that processors only make steal attempts during the execution of idle iterations, and make exactly one steal attempt for each such iteration, exactly I steal attempts take place during a computation's execution. As such, the number of CAS instructions executed in situations like this one is at most $O(I)$.

Empty is returned In this case p will not have an assigned node at the end of the scheduling iteration's execution. Thus, after p finishes executing the iteration, two scenarios may occur:

p executes an idle iteration For this case, we can create a mapping from idle iterations to each busy iteration that precedes an idle iteration, implying there can be at most $O(I)$ such iterations. With this, it is trivial to conclude that the number of CAS instructions executed by Low-Cost Work Stealing for situations equivalent to this one is at most $O(I)$.

The execution terminates Since there are exactly P processors, at most P scheduling iterations can precede the end of a computation's execution. Con-

sequently, the number of CAS instructions executed for scenarios equivalent to this one is at most $O(P)$.

Summing up all the possible scenarios, the number of CAS instructions executed by Low-Cost Work Stealing is at most $O(I + P)$.

We now turn to the number of MFence instructions executed during a computation's execution. To that end, we first bound the number of scheduling iterations that can contain MFence instructions. Consider any scheduling iteration s during a computation's execution, and let p denote the processor that executed the iteration. Iteration s was either an idle or a busy iteration.

s was an idle iteration By definition, at most I iterations are idle, implying there are $O(I)$ such iterations that could contain MFence instructions.

s was a busy iteration When p checks its *targeted* flag, one of the two following situations arises:

targeted is TRUE By observing Algorithm 1 we conclude that such a situation can only occur if another processor q has set p 's *targeted* to TRUE, which can only occur if q was executing an idle iteration. After executing the conditional statement, p resets its *targeted* flag back to FALSE. Thus, the total number of busy iterations where a processor has its flag set to *targeted* is at most I , because each such iteration can be mapped by an idle iteration. Consequently, the number of iterations similar to this one is at most $O(I)$.

targeted is FALSE As p is executing a busy iteration, it will execute the node it has assigned. From that node's execution, either 0, 1 or 2 other nodes can be enabled.

More than 0 nodes are enabled Lemma 5 implies that no MFence instruction is executed in this case.

0 nodes are enabled In this case, p cannot immediately assign a new node, because it did not enable any. By Algorithm 1, p will then invoke the POP method to its own split deque. With this, one of two possible situations arises:

split deque's private part is not empty As a consequence of Lemma 5, no MFence instruction is executed in this case.

split deque's private part is empty In this case, by observing Algorithm 2 we conclude that the invocation returns the special value RACE, implying p will make an invocation to POPBOTTOM still during that same iteration. From that invocation, two outcomes are possible:

A node is returned In this situation, p assigns the node. By observing Algorithm 1 it is trivial to conclude that this scenario only

arises if some processor previously set p 's *targeted* flag to TRUE. As a consequence, p transferred a node from the private part of its split deque to the public part. Again, using the same reasoning as for the case where p 's *targeted* flag is set to TRUE, we conclude the number of such iterations is at most $O(I)$.

EMPTY is returned After p finishes executing the current scheduling iteration s , two scenarios may occur: *p executes an idle iteration*: It is easy to deduce that we can create a mapping from idle iterations to each iteration satisfying the same conditions as s . Thus, there can be at most $O(I)$ such iterations. *The execution terminates*: As there are exactly P processors, at most P scheduling iterations can precede the end of a computation's execution. Consequently, there are at most P scheduling iteration similar to s .

Now, we sum up all the scheduling iterations that may contain MFence instructions. Accounting with all possible scenarios, it follows that at most $O(I + P)$ scheduling iterations may contain MFence instructions. Since any scheduling iteration is composed by at most C instructions, at most C MFence instructions can be executed per iteration, implying the number of MFence instructions executed during a computation's execution is at most $O(I + P)$. □

4.1 Bounds on the expected number of idle iterations

The rest of the analysis focuses on bounding the number of idle iterations that take place during a computation's execution, and follows the same general arguments as the analysis presented in Arora et al., (2001).

We say that a node u is *stealable* if u is stored in the public part of some processor's split deque. Furthermore, we denote the set of ready nodes at some step i by R_i . Consider any node $u \in R_i$. The potential associated with u at step i is denoted by $\phi_i(u)$ and is defined as

$$\phi_i(u) = \begin{cases} 4^{3w(u)-2} & \text{if } u \text{ is assigned} \\ 4^{3w(u)-1} & \text{if } u \text{ is stealable} \\ 4^{3w(u)} & \text{otherwise.} \end{cases}$$

The total potential at step i , denoted by Φ_i , corresponds to the sum of potentials of all the nodes that are ready at that step: $\Phi_i = \sum_{u \in R_i} \phi_i(u)$.

The following lemma is a formalization of the arguments already given in Arora et al., (2001), but considering the potential function we present.

Lemma 7 Consider some node u , ready at step i during the execution of a computation.

1. If u gets assigned to a processor at that step, the potential drops by at least $\frac{3}{4}\phi_i(u)$.
2. If u becomes stealable at that step, the potential drops by at least $\frac{3}{4}\phi_i(u)$.
3. If u was already assigned to a processor and gets executed at that step i , the potential drops by at least $\frac{47}{64}\phi_i(u)$.

For the remainder of the analysis, we make use of a few more definitions, first introduced in Arora et al., (2001). We denote the set of ready nodes attached to some processor p (i.e., the ready nodes in p 's split deque together with the node it has assigned, if any) at the beginning of some step i by $R_i(p)$. Furthermore, we define the total potential associated with p at step i as the sum of the potentials of each of the nodes that is attached to p at the beginning of that step $\Phi_i(p) = \sum_{u \in R_i(p)} \phi_i(u)$.

For each step i , we partition the processors into two sets, D_i and A_i , where the first is the set of all processors whose split deque is not empty at the beginning of step i , while the second is the set of all other processors (i.e., the set of all processors whose split deque is empty at the beginning of that step). Thus, the potential of any step i , Φ_i , is composed by the potential associated with each of these two partitions $\Phi_i = \Phi_i(D_i) + \Phi_i(A_i)$, where $\Phi_i(D_i) = \sum_{p \in D_i} \Phi_i(p)$ and $\Phi_i(A_i) = \sum_{p \in A_i} \Phi_i(p)$.

The following lemma is a direct consequence of Corollary 1 and of the potential function's properties.

Lemma 8 Consider any step i and any processor $p \in D_i$. The top-most node u in p 's split deque contributes at least $\frac{4}{5}$ of the potential associated with p . That is, we have $\phi_i(u) \geq \frac{4}{5}\Phi_i(p)$.

With this, we now show that if a processor p is targeted by a steal attempt, then p 's potential decreases by a constant factor. Recall C denotes a large enough constant such that any sequence of C instructions executed by a processor includes a milestone.

Lemma 9 Suppose a thief processor p chooses a processor $q \in D_i$ as its victim at some step j , such that $j \geq i$ (i.e., a steal attempt of p targeting q occurs at step j). Then, at step $j + 2C$, the potential decreased by at least $\frac{3}{5}\Phi_i(q)$ due either to assigning the topmost node in q 's split deque, or to making the topmost node of q 's split deque become stealable.

The next lemma is a trivial generalization of the original result presented in Arora et al., (2001), Balls and Weighted Bins.

Lemma 10 (Balls and Weighted Bins) Suppose we are given at least B balls and exactly B bins. Each of the balls is

tossed independently and uniformly at random into one of the B bins, where for $i = 1, \dots, B$, bin i has a weight W_i . The total weight is $W = \sum_{i=1}^B W_i$. For each bin i , we define the random variable X_i as

$$X_i = \begin{cases} W_i & \text{if some ball lands in bin } i \\ 0 & \text{otherwise} \end{cases}$$

and define the random variable X as $X = \sum_{i=1}^B X_i$. Then, for any β in the range $0 < \beta < 1$, we have $P\{X \geq \beta W\} \geq 1 - \frac{1}{(1-\beta)e}$.

The following result states that for each P idle iterations that take place, with constant probability, the total potential drops by a constant factor. The result is a consequence of Lemmas 9 and 10. Again, recall that any sequence of instructions with length at least C includes a milestone.

Lemma 11 Consider any step i and any later step j such that at least P idle iterations occur from i (inclusive) to j (exclusive). Then, we have

$$P\left\{\Phi_i - \Phi_{j+2C} \geq \frac{3}{10}\Phi_i(D_i)\right\} > \frac{1}{4}.$$

Following Lemma 11, we are able to bound the expected number of idle iterations that take place during a computation's execution using the Low-Cost Work Stealing algorithm.

Lemma 12 Consider any computation with work T_1 and critical-path length T_∞ being executed by Low-Cost Work Stealing using P processors. The expected number of idle iterations is at most $O(PT_\infty)$, and with probability at least $1 - \varepsilon$, the number of idle iterations is at most $O\left((T_\infty + \ln\left(\frac{1}{\varepsilon}\right))P\right)$.

Finally, using Lemma 12, we can obtain bounds on both expected runtime of computations executed by the Low-Cost Work Stealing algorithm, and the associated synchronization overheads.

Theorem 1 Consider any computation with work T_1 and critical-path length T_∞ being executed by the Low-Cost Work Stealing algorithm with P processors. The expected execution time is at most $O\left(\frac{T_1}{P} + T_\infty\right)$, and with probability at least $1 - \varepsilon$, the execution time is at most $O\left(\frac{T_1}{P} + T_\infty + \ln\left(\frac{1}{\varepsilon}\right)\right)$. Moreover, the expected number of CAS and MFence instructions executed during the computation's execution caused by Low-Cost Work Stealing is at most $O(PT_\infty)$, and with probability at least $1 - \varepsilon$ the number of CAS and MFence instructions executed is at most $O\left(P\left(T_\infty + \ln\left(\frac{1}{\varepsilon}\right)\right)\right)$.

Proof Both results follow directly from Lemmas 4, 6 and 12. \square

Corollary 2 Consider the statement of Theorem 1. Furthermore, let C_{CAS} and C_{MFence} denote, respectively, the maximum synchronization overheads incurred by the execution of a CAS and MFence instructions. The expected synchronization overheads incurred by Low-Cost Work Stealing are at most $O\left((C_{\text{CAS}} + C_{\text{MFence}})PT_\infty\right)$, and with probability at least $1 - \varepsilon$ the synchronization overheads incurred by Low-Cost Work Stealing are at most $O\left((C_{\text{CAS}} + C_{\text{MFence}})P\left(T_\infty + \ln\left(\frac{1}{\varepsilon}\right)\right)\right)$.

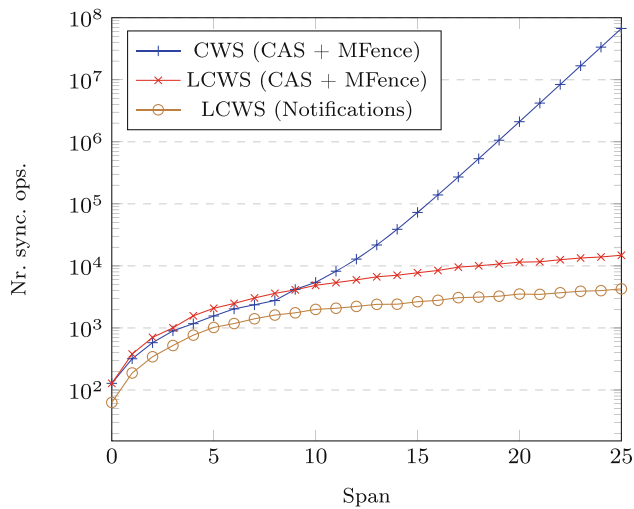
Proof As already mentioned, and by considering the definition of Low-Cost Work Stealing, depicted in Algorithm 1, the only synchronization mechanisms the scheduler uses are CAS and MFence instructions. This corollary is then a direct consequence of Theorem 1 that takes into account the maximum possible overhead incurred by the execution of a single CAS and MFence instructions. \square

5 Comparison with work stealing

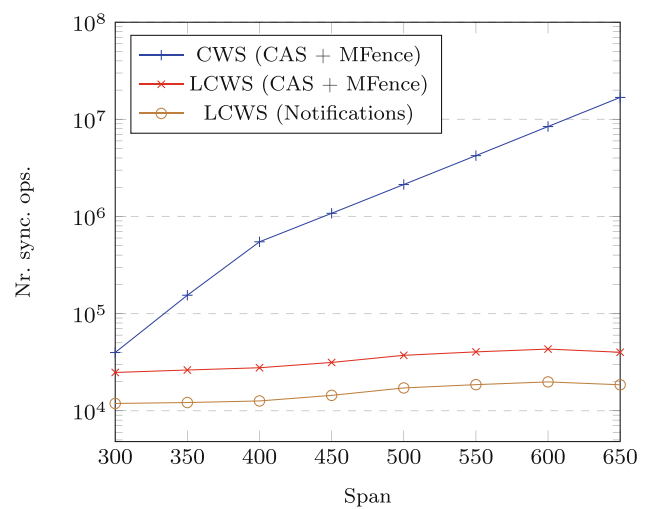
To get a better understanding of the importance of avoiding synchronization for local deque accesses, we now compare the synchronization costs of our algorithm against conventional Work Stealing algorithms that use concurrent dequeues. To that end, we developed a simulator that, given a computation's dag, executes it, monitoring not only the number of CAS and MFence instructions executed but also the number of times that thieves requested other processors to expose work. In this section, we use the term *notification* to refer to when a thief sets another processor's *targeted* flag to TRUE, requesting it to expose work.

For this comparison, we consider two distinct classes of dags: *regular* and *irregular*. Regular dags essentially correspond to trees of instructions where every non-leaf instruction forks two other instructions, and whose depth is given by an argument that is passed to the simulator. Irregular Dags are intended to simulate unbalanced computations. To that end, we use the argument passed to the simulator as the total depth of the dag and make the depth between each two consecutive fork instructions follow an exponential distribution with parameter $\lambda = 0.05$. The first class of dags corresponds to computations with balanced parallelism (e.g., Fibonacci, Parallel-For, etc.) whilst the second corresponds to the ones with unbalanced parallelism (e.g., Graph Searches).

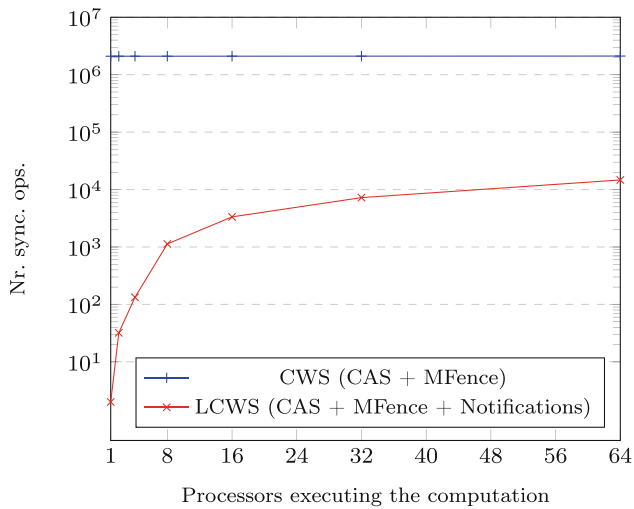
From Fig. 2a, it is clear that while for Work Stealing with concurrent dequeues the number of synchronization operations grows linearly with the total amount of work (and thus exponentially increases with the span of the dag), for Low-Cost Work Stealing the number of synchronization operations and notifications scales linearly with the span of the computation. Thus, even if the costs of handling notifications (i.e., of



(a) 64 procs, #sync ops vs span, regular dags.



(b) 64 procs, #sync ops vs span, irregular dags.



(c) #sync ops vs #procs, regular dags.

Fig. 2 Comparison of Low-Cost Work Stealing (LCWS) with Classical Work Stealing (CWS)

exposing work) were a thousand times greater than the cost of executing CAS or MFence instructions, for dags with fork-span of at least ≈ 20 , our algorithm would incur in less synchronization overheads. In practice, computations with a fork-span ≥ 20 are very common, especially among fine-grained parallelism. Unfortunately, due to the limitations that come with using a simulator, we have not been able to benchmark dags with a fork-span greater than 25. Yet, we remark that the trend is clear and confirms that the use of split deque allows to avoid most of the synchronization that is present in conventional Work Stealing algorithms. Figure 2b reinforces our insight, showing that even for computations exhibiting irregular parallelism, Low-Cost Work Stealing is able to avoid most of the synchronization present in Work Stealing algorithms that use concurrent deques. Finally, Fig. 2c shows

that while the synchronization costs for Work Stealing are always extremely high, even for single processor executions, for Low-Cost Work Stealing these costs only grow linearly with the number of processors used and, for a single processor execution, synchronization is negligible.

From a more theoretical perspective, note that, by taking into account our assumptions (which are standard: Acar et al., 2002, 2013; Agrawal et al., 2007, 2008; Arora et al., 1998, 2001; Blumofe & Leiserson, 1999; Muller & Acar, 2016; Tchiboukdjian et al., 2010) regarding computations' structure, we can create computations for which $T_1 = O(2^{T_\infty})$ (which correspond to dags of the first class). Since for such computations the number of deque accesses is directly proportional to the total amount of work (T_1), our result shows that the use of split deques allows one to reduce by almost

an exponential factor the synchronization present in conventional Work Stealing algorithms.

6 Related work

Many efforts have been carried out toward reducing and even eliminating the expensive synchronization present in state-of-the-art Work Stealing schedulers.

Michael et al. proposed an idempotent version of Work Stealing that eliminates the overheads of synchronization by relaxing the semantics of work queues (Michael et al., 2009). Concretely, rather than using the conventional *exactly once* semantics, the authors present several concurrent data structures only satisfying *at-least-once* semantics. By using these structures, processors no longer have to incur expensive synchronization overheads when operating locally, which, as the authors show, can be extremely beneficial. Unfortunately this strategy suffers from two limitations: On the one hand, for non-idempotent computations the synchronization overheads are moved to the computation itself, as it is necessary to ensure correctness; on the other, for idempotent computations these semantics can lead to situations where some computationally heavy tasks are executed more than once, limiting the scheduler's performance.

Endo et al. were the first to use split queues to avoid unnecessary synchronization (Endo et al., 1997). In this study, the authors present an implementation of a scalable garbage collector system that, by using clever load balancing techniques and split queues to avoid unnecessary synchronization overheads, achieves high performance even for large-scale machines.

In (Dinan et al., 2008, 2009), Work Stealing is studied under a distributed environment and the use of split dequeues to avoid synchronization for local deque accesses is proposed. The authors showcased the practical advantages of using split dequeues in large scale distributed settings, comparing the performance of their algorithm when using split dequeues versus when using concurrent dequeues.

Lifflander et al. studied the execution of iterative over-decomposed applications (Lifflander et al., 2012) and proposed, among others, a message-based retentive Work Stealing algorithm adapted for the execution of iterative workloads on large scale distributed settings. To avoid synchronization overheads and improve the overall performance of the scheduler, their Work Stealing algorithm uses split dequeues. Their evaluation shows that Work Stealing (with split dequeues) is a practical algorithm even for systems with hundreds of thousands of processors. A distinction between the Work Stealing algorithm proposed by Lifflander et al. and Low-Cost Work Stealing is that our algorithm keeps processors' dequeues entirely private by default. As we will discuss in

Sect. 7, this fact is key for guaranteeing that a 1-Processor execution of our algorithm requires no synchronization.

Tzannes et al. proposed a scheduling algorithm where each processor keeps all of its work entirely private, except for the topmost node that is kept stored in a shared cell (Tzannes et al., 2011). Since the algorithm always ensures that the topmost node is shared, it does not behave appropriately for computations in which processors frequently access the topmost nodes of their dequeues. As mentioned in Acar et al., (2013), a similar limitation has been identified for the Chase-Lev Deque (Chase & Lev, 2005). Unfortunately, in all these approaches (Dinan et al., 2008, 2009; Lifflander et al., 2012; Tzannes et al., 2011), processors expose work too eagerly, always leaving some work exposed for thieves to take. A consequence of this design choice is that synchronization overheads still scale with the total amount of work.

Hiraishi et al. suggested that processors should behave as in a sequential execution, thus eliminating all overheads inherent from parallelism (Hiraishi et al., 2009). Under their scheme, dequeues are kept entirely private and processors only permit parallelism when an idle processor requests work. Upon such request, the busy processor backtracks to the last point where it could have spawned a task, spawns the task, offers it to the requesting processor, and, finally, proceeds with the execution. Since work requests are usually rare, the gains of eliminating synchronization for local operation greatly surpass the extra overheads caused by using entirely private dequeues, which for this algorithm come from backtracking. This remark is key and motivated the efforts for reducing synchronization overheads to focus on the elimination of synchronization for local operation.

Morrison et al. studied alternative designs to the synchronization protocols used by Work Stealing schedulers, considering the architectures of modern *TSO* processors (Morrison et al., 2014). The authors found that by taking into account the actual implementation of microprocessors, MFence instructions can be fully eliminated while maintaining correctness. In their algorithm, thieves can only steal work from a victim if such work is stored at a safe distance, far enough from the bottom of the victim's deque to avoid any data race. This safe distance is computed *a priori*, taking into account the size of the microprocessor's internal store buffer to determine the minimum safe distance to avoid any possible conflicts. With this strategy, not only thieves can asynchronously take work from their victims, but processors can also access their dequeues locally without requiring any synchronization. Unfortunately, their scheme suffers from a limitation similar to Dijk et al.'s, as the bottommost threads within a processor's deque cannot be stolen, meaning their strategy is not appropriate for generic computations (van Dijk & van de Pol, 2014; van Ede, 2015). Nevertheless, the authors' showed that this technique (similar to using split

deque) substantially outperforms Work Stealing algorithms that use concurrent dequeues.

More recently, Dijk et al. studied the effectiveness of split dequeues on shared memory environments (van Dijk & van de Pol, 2014; van Ede, 2015). In their approach, however, busy processors only check for work requests each time they access their deque, which precludes any performance guarantees for generic computations. This is since the frequency at which busy processors may permit load balancing depends on the structure of the computation. In particular, for computations whose workload is mostly composed by large sequential threads, processors will rarely check for the necessity of load balancing, giving a very small room for parallelism. Moreover, whenever a busy processor realizes it was targeted by a steal attempt, it exposes at least half of its work. This strategy increases the unnecessary synchronization costs incurred by the algorithm as busy processors now have to start accessing the shared part of their work-queue more often to fetch work. In conclusion, although this work builds from an idea similar to ours, the proposed design choices preclude performance guarantees for generic computations.

Acar et al. presents two Work Stealing algorithms—*sender-* and *receiver-initiated*—that avoid synchronization by making dequeues entirely private to each processor (Acar et al., 2013). In addition to promising empirical results, the authors show that the expected execution time for both algorithms can be somewhat competitive with Work Stealing algorithms that use concurrent dequeues. Unfortunately, for the sender-initiated algorithm, busy processors now have to periodically search for idle ones, leading to unnecessary communication and synchronization overheads that still scale with the computation's execution time, and thus, indirectly, with the total amount of work. The difference between the receiver-initiated algorithm and ours is more subtle, however. In their receiver-initiated algorithm, busy processors now have to periodically check for incoming steal requests as well as to expose part of their current state by means of a flag that indicates whether they have any work to offer. By carefully analyzing Acar et al.'s receiver-initiated algorithm, one can see that indeed for checking and updating such exposed state, a memory fence is required.¹ This contrasts with our work, which does not require any sort of instruction synchronization. More concretely, our algorithm only requires reading a fresh Boolean flag from the memory, an operation not requiring any type of synchronization. This difference is reflected, for example, in a sequential execution: While

our algorithm essentially does not use synchronization, the receiver-initiated algorithm does.

7 Conclusion

In this paper, we studied a Work Stealing algorithm that uses split dequeues to reduce synchronization overheads. Whereas traditional Work Stealing algorithms require synchronization for every time processors access dequeues, in our proposal, synchronization operations are employed optimally, which is the key for eliminating most unnecessary synchronization overheads. By default, busy processors operate locally on their dequeues without any synchronization, resembling a sequential execution. Idle processors can request busy ones to expose some of their work, thus allowing for load balancing via direct steals. This lazy approach for using synchronization is the key for guaranteeing an asymptotically optimal expected runtime while provably reducing synchronization overheads. Indeed, we proved that the expected total synchronization of the algorithm is $O(PT_\infty(C_{CAS} + C_{MFence}))$. To justify the tightness of our bounds, we recall that, for Low-Cost Work Stealing, the expected number of (successful and unsuccessful) steal attempts is $O(PT_\infty)$. By noting that the public part of a split deque is essentially a concurrent deque, and, by taking into account the impossibility of eliminating all synchronization from the implementation of concurrent dequeues while maintaining their correctness (see Attiya et al., 2011), we conclude that the synchronization bounds we have obtained for Low-Cost Work Stealing are tight.

Finally, and as already discussed in Sect. 5, for several types of computations, the synchronization overheads of conventional Work Stealing algorithms grow linearly with both the total amount of work and the number of steal attempts. For numerous classes of parallel computations, the total amount of work increases exponentially with the span of the computation (i.e., $T_1 = O(2^{T_\infty})$). From this perspective, our results make evident the significance of the improvement of Low-Cost Work Stealing over prior Work Stealing algorithms: Not only are the synchronization overheads incurred by our algorithm (essentially) exponentially smaller than previous algorithms, but our algorithm also maintains the asymptotically optimal expected runtime bounds of the concurrent deque Work Stealing algorithm (Arora et al., 2001).

Acknowledgements This work is supported by NOVA LINCS c(UIDB/04516/2020) with the financial support of FCT—Fundação para a Ciência e a Tecnologia, through national funds. We would also like to thank Tomas Hruz and the anonymous reviewers for the helpful feedback.

¹ In particular, a processor p_i must execute an MFence instruction after writing the variable $a[i]$ (in the `update_status` method) to guarantee that idle processors learn, in constant time, that p_i has work to be stolen (see Sewell et al., 2010), as is required to achieve the expected runtime bounds presented in their paper.

A Proofs

Lemma 3 is crucial for the performance analysis of Low-Cost Work Stealing. An analogous result has already been proved for concurrent dequeues (see Arora et al., 2001, Lemma 3). For the sake of completeness we present its proof, which is a simple transcription of original proof of Arora et al. (2001, Lemma 3), adapted for split dequeues.

Lemma 3 (Structural Lemma for split dequeues) *Let v_1, \dots, v_k denote the nodes stored in some processor p 's split deque, ordered from the bottom of the split deque to the top, at some point in the linearized execution of Low-Cost Work Stealing. Moreover, let v_0 denote p 's assigned node (if any), and for $i = 0, \dots, k$ let u_i denote the designated parent of v_i in the enabling tree. Then, for $i = 1, \dots, k$, u_i is an ancestor of u_{i-1} in the enabling tree, and although v_0 and v_1 may have the same designated parent (i.e., $u_0 = u_1$), for $i = 2, 3, \dots, k$, $u_{i-1} \neq u_i$ (i.e., the ancestor relationship is proper).*

Proof Fix a particular split deque. The split deque state and assigned node only change when the assigned node is executed or a thief performs a successful steal. We prove the claim by induction on the number of assigned-node executions and steals since the split deque was last empty. In the base case, if the split deque is empty, then the claim holds vacuously. We now assume that the claim holds before a given assigned-node execution or successful steal, and we will show that it holds after. Specifically, before the assigned-node execution or successful steal, let v_0 denote the assigned node; let k denote the number of nodes in the split deque; let v_1, \dots, v_k denote the nodes in the split deque ordered from the bottom to top; and for $i = 0, \dots, k$, let u_i denote the designated parent of v_i . We assume that either $k = 0$, or for $i = 1, \dots, k$, node u_i is an ancestor of u_{i-1} in the enabling tree, with the ancestor relationship being proper, except possibly for the case $i = 1$. After the assigned-node execution or successful steal, let v_0' denote the assigned node; let k' denote the number of nodes in the split deque; let $v_1', \dots, v_{k'}$ denote the nodes in the split deque ordered from bottom to top; and for $i = 1, \dots, k'$, let u_i' denote the designated parent of v_i' . We now show that either $k' = 0$, or for $i = 1, \dots, k'$, node u_i' is an ancestor of u_{i-1}' in the enabling tree, with the ancestor relationship being proper, except possibly for the case $i = 1$.

Consider the execution of the assigned node v_0 by the owner.

If the execution of v_0 enables 0 children, then the owner pops the bottommost node off its split deque and makes that node its new assigned node. If $k = 0$, then the split deque is empty; the owner does not get a new assigned node; and $k' = 0$. If $k > 0$, then the bottommost node v_1 is popped and becomes the new assigned node, and $k' = k - 1$. If $k = 1$, then $k' = 0$. Otherwise, $k' = k - 1$. We now rename the

nodes as follows. For $i = 0, \dots, k'$, we set $v_i' = v_{i+1}$ and $u_i' = u_{i+1}$. We now observe that for $i = 1, \dots, k'$, node u_i' is a proper ancestor of u_{i-1}' in the enabling tree.

If the execution of v_0 enables 1 child x , then x becomes the new assigned node; the designated parent of x is v_0 ; and $k' = k$. If $k = 0$, then $k' = 0$. Otherwise, we can rename the nodes as follows. We set $v_0' = x$; we set $u_0' = v_0$; and for $i = 1, \dots, k'$, we set $v_i' = v_i$ and $u_i' = u_i$. We now observe that for $i = 1, \dots, k'$, node u_i' is a proper ancestor of u_{i-1}' in the enabling tree. That u_1' is a proper ancestor of u_0' in the enabling tree follows from the fact that (u_0, v_0) is an enabling edge.

In the most interesting case, the execution of the assigned node v_0 enables 2 children x and y , with x being pushed onto the bottom of the split deque and y becoming the new assigned node. In this case, (v_0, x) and (v_0, y) are both enabling edges, and $k' = k + 1$. We now rename the nodes as follows. We set $v_0' = y$; we set $u_0' = v_0$; we set $v_1' = x$; we set $u_1' = v_0$; and for $i = 2, \dots, k'$, we set $v_i' = v_{i-1}$ and $u_i' = u_{i-1}$. We now observe that $u_1' = u_0'$, and for $i = 2, \dots, k'$, node u_i' is a proper ancestor of u_{i-1}' in the enabling tree. That u_2' is a proper ancestor of u_1' in the enabling tree follows from the fact that (u_0, v_0) is an enabling edge.

Finally, we consider a successful steal by a thief. In this case, the thief pops the topmost node v_k off the split deque, so $k' = k - 1$. If $k = 1$, then $k' = 0$. Otherwise, we can rename the nodes as follows. For $i = 0, \dots, k'$, we set $v_i' = v_i$ and $u_i' = u_i$. We now observe that for $i = 1, \dots, k'$, node u_i' is an ancestor of u_{i-1}' in the enabling tree, with the ancestor relationship being proper, except possibly for the case $i = 1$. \square

Corollary 1 *If v_0, v_1, \dots, v_k are as defined in the statement of Lemma 3, then we have $w(v_0) \leq w(v_1) < \dots < w(v_{k-1}) < w(v_k)$.*

We are now able to bound the execution time of a computation depending on the number of idle iterations that take place during that computation's execution. The following result is a trivial variant of Arora et al. (2001, Lemma 5) but considering the Low-Cost Work Stealing algorithm, and is only added for the sake of completion.

Lemma 4 *Consider any computation with work T_1 being executed by P processors, under Low-Cost Work Stealing. The execution time is $O\left(\frac{T_1}{P} + \frac{I}{P}\right)$, where I denotes the number of idle iterations executed by processors.*

Proof Consider two buckets to which we add tokens during the computation's execution: the *busy* bucket and the *idle* bucket. At the end of each iteration, every processor places a token into one of these buckets. If a processor executed a node during the iteration, it places a token into the busy

bucket, and otherwise, it places a token into the idle bucket. Since we have P processors, for each C consecutive steps, at least P tokens are placed into the buckets.

Because, by definition, the computation has T_1 nodes, there will be exactly T_1 tokens in the busy bucket when the computation's execution ends. Moreover, as I denotes the number of idle iterations, it also corresponds to the number of tokens in the idle bucket when the computation's execution ends. Thus, exactly $T_1 + I$ tokens are collected during the computation's execution. Taking into account that for each C consecutive steps at least P tokens are placed into the buckets, we conclude the number of steps required to collect all the tokens is at most $C \cdot \left(\frac{T_1}{P} + \frac{I}{P}\right)$. After collecting all the T_1 tokens, the computation's execution terminates, implying the execution time is at most $O\left(\frac{T_1}{P} + \frac{I}{P}\right)$. \square

The following lemma is a formalization of the arguments already given in Arora et al., (2001), but considering the potential function we present.

Lemma 7 Consider some node u , ready at step i during the execution of a computation.

1. If u gets assigned to a processor at that step, the potential drops by at least $\frac{3}{4}\phi_i(u)$.
2. If u becomes stealable at that step, the potential drops by at least $\frac{3}{4}\phi_i(u)$.
3. If u was already assigned to a processor and gets executed at that step i , the potential drops by at least $\frac{47}{64}\phi_i(u)$.

Proof Regarding the first claim, if u was stealable the potential decreases from $4^{3w(u)-1}$ to $4^{3w(u)-2}$. Otherwise, the potential decreases from $4^{3w(u)}$ to $4^{3w(u)-2}$, which is even more than in the previous case. Given that $4^{3w(u)-1} - 4^{3w(u)-2} = \frac{3}{4}\phi_i(u)$, we conclude that if u gets assigned the potential decreases by at least $\frac{3}{4}\phi_i(u)$.

Regarding the second one, note that u was not stealable (because it became stealable at step i) and so the potential decreases from $4^{3w(u)}$ to $4^{3w(u)-1}$. So, if u becomes stealable, the potential decreases by $4^{3w(u)} - 4^{3w(u)-1} = \frac{3}{4}\phi_i(u)$.

We now prove the last claim. Recall that, by our conventions regarding computations' structure, each node within a computation's dag can have an out-degree of at most two. Consequently, each node can be the designated parent of at most two other ones in the enabling tree. Moreover, by definition, the weight of any node is strictly smaller than the weight of its designated parent, since it is deeper in the enabling tree than its designated parent. Consider the three possible scenarios:

0 nodes enabled The potential decreased by $\phi_i(u)$.

1 node enabled The enabled node becomes the assigned node of the processor (that executed u). Let x denote the

enabled node. Since x is the child of u in the enabling tree, it follows $\phi_i(u) - \phi_{i+1}(x) = 4^{3w(u)-2} - 4^{3w(x)-2} = 4^{3w(u)-2} - 4^{3(w(u)-1)-2} = \frac{63}{64}\phi_i(u)$. Thus, for this situation, the potential decreases by $\frac{63}{64}\phi_i(u)$.

2 nodes enabled In this case, one of the enabled nodes immediately becomes the assigned node of the processor whilst the other is pushed onto the bottom of the split deque's private part. Let x denote the enabled node that becomes the processor's new assigned node and y the other enabled node. Since both x and y have u as their designated parent in the enabling tree, we have $\phi_i(u) - \phi_{i+1}(x) - \phi_{i+1}(y) = 4^{3w(u)-2} - 4^{3w(x)-2} - 4^{3w(y)-2} = \frac{47}{64}\phi_i(u)$. As such, the potential decreases by $\frac{47}{64}\phi_i(u)$, concluding the proof of the lemma. \square

The following lemma is a direct consequence of Corollary 1 and of the potential function's properties. The result is a variant of Arora et al., (2001), Top-Heavy Deques, considering split deques instead of the conventional fully concurrent deques, and our potential function, instead of the original.

Lemma 8 Consider any step i and any processor $p \in D_i$. The top-most node u in p 's split deque contributes at least $\frac{4}{5}$ of the potential associated with p . That is, we have $\phi_i(u) \geq \frac{4}{5}\Phi_i(p)$.

Proof This lemma follows from Corollary 1. We prove it by induction on the number of nodes within p 's split deque.

Base case As the base case, consider that p 's split deque contains a single node u . The processor itself can either have or not an assigned node. For the second scenario, we have $\phi_i(u) = \Phi_i(p)$. Regarding the first case, let x denote p 's assigned node. Corollary 1 implies that $w(u) \geq w(x)$. It follows $\Phi_i(p) = \phi_i(u) + \phi_i(x) = 4^{3w(u)-1} + 4^{3w(x)-2} \leq \frac{5}{4}\phi_i(u)$. Thus, if p 's split deque contains a single node we have $\Phi_i(p) \leq \frac{5}{4}\phi_i(u)$.

Induction step Consider that p 's split deque now contains n nodes, where $n \geq 2$, and let u, x denote the topmost and second topmost nodes, respectively, within the split deque. For the purpose of induction, let us assume the lemma holds for all the first $n - 1$ nodes (i.e., without accounting with u): $\Phi_i(p) - \phi_i(u) \leq \frac{5}{4}\phi_i(x)$. Corollary 1 implies $w(u) > w(x) \equiv w(u) - 1 \geq w(x)$. It follows $\Phi_i(p) \leq \frac{5}{4}\phi_i(x) + \phi_i(u) = \frac{5}{4}4^{3w(x)} + 4^{3w(u)-1} \leq \frac{5}{4}4^{3(w(u)-1)} + 4^{3w(u)-1} = \frac{69}{64}\phi_i(u) < \frac{5}{4}\phi_i(u)$ concluding the proof of the lemma. \square

The following result is a consequence of Lemma 8.

Lemma 9 Suppose a thief processor p chooses a processor $q \in D_i$ as its victim at some step j , such that $j \geq i$ (i.e., a steal attempt of p targeting q occurs at step j). Then, at step $j + 2C$, the potential decreased by at least $\frac{3}{5}\Phi_i(q)$ due to

either the assignment of the topmost node in q 's split deque, or for making the topmost node of q 's split deque become stealable.

Proof Let u denote the topmost node of q 's split deque at the beginning of step i . We first prove that u either gets assigned or becomes stealable.

Three possible scenarios may take place due to p 's steal attempt targeting q 's split deque.

The invocation returns a node If p stole u , then, u gets assigned to p . Otherwise, some other processor removed u before p did, implying u got assigned to that other processor.

The invocation aborts Since the split deque implementation meets the relaxed semantics on any good set of invocations, and because the Low-Cost Work Stealing algorithm only makes good sets of invocations, we conclude that some other processor successfully removed a topmost node from q 's split deque during the aborted steal attempt made by p . If the removed node was u , u gets assigned to a processor (that may either be q , or, some other thief that successfully stole u). Otherwise, u must have been previously stolen by a thief or popped by q , and thus became assigned to some processor.

The invocation returns EMPTY This situation can only occur if either q 's split deque is completely empty, or if there is no node in the public part of q 's split deque.

- For the first case, since $q \in D_i$, some processor must have successfully removed u from q 's split deque. Consequently, u was assigned to a processor.
- If there was no node in the public part of q 's split deque, p sets q 's *targeted* flag to TRUE in a later step j' . Recall that, for each C consecutive instructions executed by a processor, at least one corresponds to a milestone. It follows that $j' \leq j + C$. Furthermore, by observing Algorithm 1, we conclude that q will make and complete an invocation to UPDATEBOTTOM of its split deque in one of the C steps succeeding step j' . Thus, if q 's split deque's private part is not empty, a node will become stealable. From that invocation, only two possible situations can take place:

No node becomes stealable In this case, the private part of q 's split deque was empty, implying some processor (either q or some thief) assigned u .

A node becomes stealable If the node that became stealable as the result of the invocation was not u , then either u was assigned by a processor (that could have been q or some thief), or u had already been transferred to the public part of q 's split deque as a consequence of another thief's steal attempt that also returned EMPTY, implying that either u became assigned, or it became steal-

able. Otherwise, the node that became stealable as a result of the UPDATEBOTTOM's invocation was u . Thus, in any case, u either gets assigned to a processor or becomes stealable.

With this, we conclude that u either became assigned or became stealable until step $j + 2C$.

From Lemma 8, we have $\phi_i(u) \geq \frac{4}{5}\Phi_i(q)$. Furthermore, Lemma 7 proves that if u gets assigned the potential decreases by at least $\frac{3}{4}\phi_i(u)$, and if u becomes stealable the potential also decreases by at least $\frac{3}{4}\phi_i(u)$. Because u is either assigned or becomes stealable in any case, we conclude the potential associated with q at step $j + 2C$ has decreased by at least $\frac{3}{5}\Phi_i(q)$. \square

The following lemma is trivial a generalization of the original result presented in Arora et al. (2001, Balls and Weighted Bins). The only difference between the two results is the assumption of having at least B balls, rather than exactly B balls. Its proof is only presented for the sake of completion and is (trivially) adapted from the proof of Arora et al. (2001, Balls and Weighted Bins).

Lemma 10 (Balls and Weighted Bins) *Suppose we are given at least B balls, and exactly B bins. Each of the balls is tossed independently and uniformly at random into one of the B bins, where for $i = 1, \dots, B$, bin i has a weight W_i . The total weight is $W = \sum_{i=1}^B W_i$. For each bin i , we define the random variable X_i as*

$$X_i = \begin{cases} W_i & \text{if some ball lands in bin } i \\ 0 & \text{otherwise} \end{cases}$$

and define the random variable X as $X = \sum_{i=1}^B X_i$. Then, for any β in the range $0 < \beta < 1$, we have

$$P\{X \geq \beta W\} \geq 1 - \frac{1}{(1-\beta)e}.$$

Proof Consider the random variable $W_i - X_i$ taking the value of W_i when no ball lands in bin i and 0 otherwise, and let B' denote the total number of balls that are tossed. It follows $E[W_i - X_i] = W_i(1 - \frac{1}{B})^{B'} \leq \frac{W_i}{e}$. From the linearity of expectation, we have $E[W - X] \leq \frac{W}{e}$. Markov's inequality then implies $P\{W - X > (1 - \beta)W\} = P\{X < \beta W\} \leq \frac{E[W - X]}{(1-\beta)W} \leq \frac{1}{(1-\beta)e}$. \square

The following result states that for each P idle iterations that take place, with constant probability the potential drops by a constant factor. An analogous lemma was originally presented in Arora et al. (2001, Lemma 8) for the non-blocking Work Stealing algorithm. The result is a consequence of Lemmas 9 and 10, and its proof follows the same traits as the one presented in that study.

Lemma 11 Consider any step i and any later step j such that at least P idle iterations occur from i (inclusive) to j (exclusive). Then, we have

$$P \left\{ \Phi_i - \Phi_{j+2C} \geq \frac{3}{10} \Phi_i(D_i) \right\} > \frac{1}{4}.$$

Proof By Lemma 9 we know that for each processor $p \in D_i$ that is targeted by a steal attempt, the potential drops by at least $\frac{3}{5} \Phi_i(p)$, at most $2C$ steps after being targeted.

When executing an idle iteration, a processor plays the role of a thief attempting to steal work from some victim. Thus, since P idle iterations occur from step i (inclusive) to step j (exclusive), at least P steal attempts take place during that same interval. We can think of each such steal attempt as a ball toss of Lemma 10.

For each processor p in D_i , we assign it a weight $W_p = \frac{3}{5} \Phi_i(p)$, and for each other processor p in A_i , we assign it a weight $W_p = 0$. Clearly, the weights sum to $W = \frac{3}{5} \Phi_i(D_i)$. Using $\beta = \frac{1}{2}$ in Lemma 10, it follows that with probability at least $1 - \frac{1}{(1-\beta)^e} > \frac{1}{4}$, the potential decreases by at least $\beta W = \frac{3}{10} \Phi_i(D_i)$, concluding the proof of this lemma. \square

Finally, we bound the expected number of idle iterations that take place during a computation's execution using the Low-Cost Work Stealing algorithm. The result follows from Lemma 11 and is proved using similar arguments as the ones used in the proof of Arora et al. (2001, Theorem 9). The presented proof corresponds to an adaptation of the one originally presented for the just mentioned theorem.

Lemma 12 Consider any computation with work T_1 and critical-path length T_∞ being executed by Low-Cost Work Stealing using P processors. The expected number of idle iterations is at most $O(PT_\infty)$. Moreover, with probability at least $1 - \varepsilon$ the number of idle iterations is at most $O((T_\infty + \ln(\frac{1}{\varepsilon}))P)$.

Proof To analyze the number of idle iterations, we break the execution into *phases*, each composed by $\Theta(P)$ idle iterations. Then, we prove that with constant probability, a phase leads the potential to drop by a constant factor.

A computation's execution begins when the root gets assigned to a processor. By definition, the weight of the root is T_∞ , implying the potential at the beginning of a computation's execution starts at $\Phi_0 = 4^{3T_\infty-2}$. Furthermore, it is straightforward to deduce that the potential is 0 after (and only after) a computation's execution terminates. We use these facts to bound the expected number of phases needed to decrease the potential down to 0. The first phase starts at

step $t_1 = 1$ and ends at the first step t_1' such that, at least P idle iterations took place during the interval $[t_1, t_1' - 2C]$. The second phase starts at step $t_2 = t_1' + 1$, and so on.

Consider two consecutive phases starting at steps i and j , respectively. We now prove that $P \{ \Phi_j \leq \frac{7}{10} \Phi_i \} > \frac{1}{4}$. Recall that we can partition the potential as $\Phi_i = \Phi_i(A_i) + \Phi_i(D_i)$. Since, from the beginning of each phase and until its last $2C$ steps, at least P idle iterations take place, then, by Lemma 11 it follows $P \{ \Phi_i - \Phi_j \geq \frac{3}{10} \Phi_i(D_i) \} > \frac{1}{4}$. Now, we have to prove the potential also drops by a constant fraction of $\Phi_i(A_i)$. Consider some processor $p \in A_i$:

- If p does not have an assigned node, then $\Phi_i(p) = 0$.
- Otherwise, if p has an assigned node u at step i , then, $\Phi_i(p) = \phi_i(u)$. Noting that each phase has more than C steps, then, p executes u before the next phase begins (i.e., before step j). Thus, the potential drops by at least $\frac{47}{64} \phi_i(u)$ during that phase.

Cumulatively, for each $p \in A_i$, it follows $\Phi_i - \Phi_j \geq \frac{47}{64} \Phi_i(A_i)$. Thus, no matter how Φ_i is partitioned between $\Phi_i(A_i)$ and $\Phi_i(D_i)$, we have $P \{ \Phi_i - \Phi_j \geq \frac{3}{10} \Phi_i \} > \frac{1}{4}$.

We say a phase is successful if it leads the potential to decrease by at least a $\frac{3}{10}$ fraction. So, a phase succeeds with probability at least $\frac{1}{4}$. Since the potential is an integer, and, as aforementioned, starts at $\Phi_0 = 4^{3T_\infty-2}$ and ends at 0, then there can be at most $(3T_\infty - 2) \log_{\frac{10}{7}}(4) < 12T_\infty$ successful phases. If we think of each phase as a coin toss, where the probability that we get heads is at least $\frac{1}{4}$, then the expected number of coins we have to toss to get heads $12T_\infty$ times is at most $48T_\infty$. In the same way, the expected number of phases needed to obtain $12T_\infty$ successful ones is at most $48T_\infty$. Consequently, the expected number of phases is $O(T_\infty)$. Moreover, as each phase contains $O(P)$ idle iterations, the expected number of idle iterations is $O(PT_\infty)$.

Now, suppose the execution takes $n = 48T_\infty + m$ phases. Each phase succeeds with probability greater or equal to $p = \frac{1}{4}$, meaning the expected number of successes is at least $np = 12T_\infty + \frac{m}{4}$. We now compute the probability that the number of X successes is less than $12T_\infty$. We use the Chernoff bound (Alon & Spencer, 1992), $P \{ X < np - a \} < e^{-\frac{a^2}{2np}}$ with $a = \frac{m}{4}$. It follows, $np - a = 12T_\infty$. Choosing $m = 48T_\infty + 16 \ln(\frac{1}{\varepsilon})$, we have $P \{ X < 12T_\infty \} < e^{-\frac{(\frac{m}{4})^2}{2(12T_\infty + \frac{m}{4})}} \leq e^{-\frac{m}{16}} \leq e^{-\frac{16 \ln(\frac{1}{\varepsilon})}{16}} = \varepsilon$. Thus, the probability that the execution takes $96T_\infty + 16 \ln(\frac{1}{\varepsilon})$ phases or more, is less than ε . With this we conclude that the number of idle iterations is at most $O((T_\infty + \ln(\frac{1}{\varepsilon}))P)$ with probability at least $1 - \varepsilon$. \square

References

- Acar, U. A., Blleloch, G. E., & Blumofe, R. D. (2002). The data locality of work stealing. *Theory of Computing Systems*, 35(3), 321–347. <https://doi.org/10.1007/s00224-002-1057-3>
- Acar, U. A., Charguéraud, A., & Rainey, M. (2013). Scheduling parallel programs by work stealing with private deques. In *ACM SIGPLAN symposium on principles and practice of parallel programming, PPOPP'13, February 23–27, 2013* (pp. 219–228). <https://doi.org/10.1145/2442516.2442538>
- Agrawal, K., He, Y., Hsu, W., & Leiserson, C. E. (2007). Adaptive scheduling with parallelism feedback. In *21th International parallel and distributed processing symposium (IPDPS 2007), proceedings, 26–30 March 2007* (pp. 1–7). <https://doi.org/10.1109/IPDPS.2007.370496>
- Agrawal, K., Leiserson, C. E., He, Y., & Hsu, W. (2008). Adaptive work-stealing with parallelism feedback. *ACM Transactions on Computing Systems*. <https://doi.org/10.1145/1394441.1394443>
- Alon, N., & Spencer, J. (1992). *The probabilistic method*. Wiley.
- Arora, N. S., Blumofe, R. D., & Plaxton, C. G. (1998). Thread scheduling for multiprogrammed multiprocessors. In *SPAA* (pp. 119–129). <https://doi.org/10.1145/277651.277678>
- Arora, N. S., Blumofe, R. D., & Plaxton, C. G. (2001). Thread scheduling for multiprogrammed multi processors. *Theory of Computing Systems*, 34(2), 115–144. <https://doi.org/10.1007/s00224-001-0004-z>
- Attiya, H., Guerraoui, R., Hendler, D., Kuznetsov, P., Michael, M. M., & Vechev, M. T. (2011). Laws of order: Expensive synchronization in concurrent algorithms cannot be eliminated. In *Proceedings of the 38th ACM SIGPLAN-SIGACT symposium on principles of programming languages, POPL 2011, January 26–28, 2011* (pp. 487–498). <https://doi.org/10.1145/1926385.1926442>
- Blumofe, R. D., Joerg, C. F., Kuszmaul, B. C., Leiserson, C. E., Randall, K. H., & Zhou, Y. (1996). Cilk: An efficient multithreaded runtime system. *Journal of Parallel and Distributed Computing*, 37(1), 55–69. <https://doi.org/10.1006/jpdc.1996.0107>
- Blumofe, R. D., & Leiserson, C. E. (1999). Scheduling multithreaded computations by work stealing. *Journal of the ACM*, 46(5), 720–748. <https://doi.org/10.1145/324133.324234>
- Blumofe, R. D., & Papadopoulos, D. (1998). The performance of work stealing in multiprogrammed environments. In *ACM sigmetrics performance evaluation review* (Vol. 26, pp. 266–267).
- Blumofe, R. D., Plaxton, C. G., & Ray, S. (1999). *Verification of a concurrent deque implementation*. University of Texas at Austin.
- Chase, D., & Lev, Y. (2005). Dynamic circular work-stealing deque. In *SPAA 2005: Proceedings of the 17th annual ACM symposium on parallelism in algorithms and architectures, July 18–20, 2005* (pp. 21–28). <https://doi.org/10.1145/1073970.1073974>
- Dinan, J., Krishnamoorthy, S., Larkins, D. B., Nieplocha, J., & Sadayappan, P. (2008). Scioto: A framework for global-view task parallelism. In *2008 International conference on parallel processing, ICPP 2008, September 8–12, 2008* (pp. 586–593). Retrieved from <https://doi.org/10.1109/ICPP.2008.44>
- Dinan, J., Larkins, D. B., Sadayappan, P., Krishnamoorthy, S., & Nieplocha, J. (2009). Scalable work stealing. In *Proceedings of the ACM/IEEE conference on high performance computing, SC 2009, November 14–20, 2009*. <https://doi.org/10.1145/1654059.1654113>
- Endo, T., Taura, K., & Yonezawa, A. (1997). A scalable mark-sweep garbage collector on large-scale shared-memory machines. In *Proceedings of the ACM/IEEE conference on supercomputing, SC 1997, November 15–21, 1997* (p. 48). Retrieved from <https://doi.org/10.1145/509593.509641>
- Frigo, M., Leiserson, C. E., & Randall, K. H. (1998). The implementation of the Cilk-5 multithreaded language. In *Proceedings of the ACM SIGPLAN '98 conference on programming language design and implementation (PLDI), June 17–19, 1998* (pp. 212–223). ACM. <https://doi.org/10.1145/277650.277725>
- Hiraishi, T., Yasugi, M., Umata, S., & Yuasa, T. (2009). Backtracking-based load balancing. In *Proceedings of the 14th ACM SIGPLAN symposium on principles and practice of parallel programming, PPOPP 2009, February 14–18, 2009* (pp. 55–64). Retrieved from <https://doi.org/10.1145/1504176.1504187>
- Lifflander, J., Krishnamoorthy, S., & Kalé, L. V. (2012). Work stealing and persistence-based load balancers for iterative overdecomposed applications. In *The 21st international symposium on high-performance parallel and distributed computing, HPDC'12, June 18–22, 2012* (pp. 137–148). Retrieved from <https://doi.org/10.1145/2287076.2287103>
- Michael, M. M., Vechev, M. T., & Saraswat, V. A. (2009). Idempotent work stealing. In *Proceedings of the 14th ACM SIGPLAN symposium on principles and practice of parallel programming, PPOPP 2009, February 14–18, 2009* (pp. 45–54). <https://doi.org/10.1145/1504176.1504186>
- Morrison, A., & Afek, Y. (2014). Fence-free work stealing on bounded TSO processors. In *Architectural support for programming languages and operating systems, ASPLOS '14, March 1–5, 2014* (pp. 413–426). <https://doi.org/10.1145/2541940.2541987>
- Muller, S. K., & Acar, U. A. (2016). Latency-hiding work stealing: Scheduling interacting parallel computations with work stealing. In *Proceedings of the 28th ACM symposium on parallelism in algorithms and architectures, SPAA 2016, July 11–13, 2016* (pp. 71–82). <https://doi.org/10.1145/2935764.2935793>
- Sewell, P., Sarkar, S., Owens, S., Nardelli, F. Z., & Myreen, M. O. (2010). x86-tso: A rigorous and usable programmer's model for x86 multiprocessors. *Communications of the ACM*, 53(7), 89–97. <https://doi.org/10.1145/1785414.1785443>
- Tchiboukdjian, M., Gast, N., Trystram, D., Roch, J., & Bernard, J. (2010). A tighter analysis of work stealing. In *Algorithms and computation—21st international symposium, ISAAC 2010, December 15–17, 2010, proceedings, part II* (pp. 291–302). https://doi.org/10.1007/978-3-642-17514-5_25
- Tzannes, A., Barua, R., & Vishkin, U. (2011). Improving run-time scheduling for general-purpose parallel code. In L. Rauchwerger & V. Sarkar (Eds.), *2011 International conference on parallel architectures and compilation techniques, PACT 2011, October 10–14, 2011* (p. 216). IEEE Computer Society. <https://doi.org/10.1109/PACT.2011.49>
- van Dijk, T., & van de Pol, J. C. (2014). Lace: Nonblocking split deque for work-stealing. In *Europar 2014: Parallel processing workshops—Europar 2014 international workshops, August 25–26, 2014, revised selected papers, part II* (pp. 206–217). https://doi.org/10.1007/978-3-319-14313-2_18
- van Ede, T. (2015). *Certainty in lockless concurrent algorithms: An informal proof of lace* (Technical Report). University of Twente.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.