

Technical Report no. 2008-69

NB-FEB: An Easy-to-Use and Scalable Universal Synchronization Primitive for Parallel Programming

Phuong Hoai Ha

Philippas Tsigas¹

Otto J. Anshus



Department of Computing Science
Faculty of Science
University of Tromsø
N-9037 Tromsø, Norway

Tromsø, October 2008.

¹Department of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Göteborg, Sweden.

Technical Report in Computing Science at
University of Tromsø

Technical Report no. 2008-69
ISSN: XXXX-XXXX

Department of Computing Science
Faculty of Science
University of Tromsø
N-9037 Tromsø, Norway

Tromsø, Norway, October 2008.

Abstract

This paper addresses the problem of universal synchronization primitives that can support scalable thread synchronization for large-scale many-core architectures. The universal synchronization primitives that have been deployed widely in conventional architectures, are the compare-and-swap (CAS) and load-linked/store-conditional (LL/SC) primitives. However, such synchronization primitives are expected to reach their scalability limits in the evolution to many-core architectures with thousands of cores.

We introduce a non-blocking full/empty bit primitive, or NB-FEB for short, as a promising synchronization primitive for parallel programming on many-core architectures. We show that the NB-FEB primitive is universal, scalable, feasible and convenient to use. NB-FEB, together with registers, can solve the consensus problem for an arbitrary number of processes (universality). NB-FEB is combinable, namely its memory requests to the same memory location can be combined into only one memory request, which consequently mitigates performance degradation due to synchronization "hot spots" (scalability). Since NB-FEB is a variant of the original full/empty bit that always returns a value instead of waiting for a conditional flag, it is as feasible as the original full/empty bit, which has been implemented in many computer systems (feasibility). The original full/empty bit is well-known as a special-purpose primitive for fast producer-consumer synchronization and has been used extensively in the specific domain of applications. In this paper, we show that NB-FEB can be deployed easily as a general-purpose primitive. Using NB-FEB, we construct a non-blocking software transactional memory system called NBFEB-STM, which can be used to handle concurrent threads conveniently. NBFEB-STM is space efficient: the space complexity of each object updated by N concurrent threads/transactions is $\Theta(N)$, the optimal.

Keywords: many-core architectures, non-blocking synchronization, full/empty bit, universal, combining, non-blocking software transactional memory, synchronization primitives.

1 Introduction

Universal synchronization primitives [28] are essential for constructing non-blocking synchronization mechanisms for parallel programming, like non-blocking software transactional memory [21, 27, 30, 36, 43]. Non-blocking synchronization eliminates the concurrency control problems of mutual exclusion locks, such as priority inversion, deadlock and convoying. As many-core architectures with thousands of cores are expected to be our future chip architectures [5], universal synchronization primitives that can support scalable thread synchronization for such large-scale architectures are desired.

However, the conventional universal primitives like *compare-and-swap* (CAS) and *load-linked/store-conditional* (LL/SC) are expected to reach their scalability limits in the evolution

to many-core architectures with thousands of cores. For each shared memory location, the LL/SC implementation conceptually associates a reservation bit with each processor. The reservations are invalidated when the location are modified by any processor. Implementing LL/SC in the memory (without compromising its semantics) limits the scalability of the multiprocessor since the total directory size increases quadratically with the number of processors [37]. Therefore, the LL/SC primitives are built on conventional cache-coherent protocols [37, 14]. However, experimental studies have shown that the LL/SC primitives are not scalable for multicore architectures [48]. The conventional cache-coherent protocols are considered inefficient for large scale manycore architectures [5]. As a result, several emerging multicore architectures like the NVIDIA CUDA [39], the ClearSpeed CSX [49], the IBM Cell BE [23] and the Cyclops-64 [12] architectures utilize fast local memory for each processing core rather than coherent data cache.

For the emerging many-core architectures without coherent data cache, the CAS primitive is not scalable either since CAS is not *combinable* [32, 10]. Primitives are combinable if their memory requests to the same memory location (arriving at a switch of the processor-to-memory interconnection network) can be combined into only one memory request. Separate replies to the original requests are later created from the reply to the combined request (at the switch). The combining technique has been implemented in the NYU Ultracomputer [22] and the IBM RP3 [41] machine and has been shown to be a promising technique for large-scale multiprocessors to alleviate the performance degradation due to synchronization "hot spot". Although the *single-valued* $CAS_a(x, b)$ [10], which will atomically swap b to x if x equals a is combinable, the number of instructions CAS_a must be as many as the number of integers a that can be stored in one memory word (e.g. 2^{64} CAS_a instructions for 64-bit words). This fact makes the *single-valued* CAS_a unfeasible for hardware implementation.

Another universal primitive called *sticky bit* has been suggested in [42], but it has not been deployed so far due to its usage complexity. To the best of our knowledge, the universal construction using the sticky bit [42] does not prevent a delayed thread, even after being helped, from jamming the sticky bits of a cell that has been re-initialized and reused. Since the universal construction is built on a doubly-linked list of cells, it is not obvious how an external garbage collector (supported by the underlying system) can help solve the problem. Moreover, the space complexity of the universal construction for an object is as high as $O(N^2 \log N)$, where N is the number of processes.

This paper suggests a novel synchronization primitive, called NB-FEB, as a promising synchronization primitive for parallel programming on many-core architectures. What makes NB-FEB be a promising primitive is its following four main properties. NB-FEB is:

Feasible : NB-FEB is a *non-blocking* variant of the conven-

tional full/empty bit that *always returns* the old value of the variable instead of waiting for its conditional flag to be set (or cleared). This simple modification makes NB-FEB as *feasible* as the original (blocking) full/empty bit, which has been implemented in many computer systems like HEP [45], Tera [3], MDP [15], Sparcle [2], M-Machine [31] and Eldorado [20]. The space overhead of full/empty bits can be reduced using the *synchronization state buffer (SSB)* [51].

Universal : This simple modification, however, significantly increases the synchronization power of full/empty bits, making NB-FEB as powerful as *CAS* or *LL/SC*. NB-FEB, together with registers, can solve consensus problem for arbitrary number of processes, the essential property for constructing non-blocking synchronization mechanisms (cf. Section 3.1).

Scalable : Like the original full/empty bit, NB-FEB is *combinable*: its memory requests to the same memory location can be combined into only one memory request (cf. Section 3.2). This empowers NB-FEB with the ability to provide *scalable* thread synchronization for large-scale many-core architectures.

Convenient to use : The original full/empty bit is well-known as a *special-purpose* primitive for fast producer-consumer synchronization and has been used extensively in the specific domain of applications. In this paper, we show that NB-FEB can be deployed easily as a *general-purpose* primitive. Using NB-FEB, we construct a non-blocking software transactional memory system called NBFEB-STM, which can be used to handle concurrent threads *conveniently*. NBFEB-STM is space efficient: the space complexity of each object updated by N concurrent threads/transactions is $\Theta(N)$, the optimal (cf. Section 4).

The rest of this paper is organized as follows. Section 2 presents the shared memory and interconnection network models assumed in this paper. Section 3 describes the NB-FEB primitive in detail and proves its universality and combinability properties. Section 4 presents NBFEB-STM, the obstruction-free multi-versioning STM constructed on the NB-FEB primitive. Section 5 describes a garbage collector that can be used as an external garbage collector for the NBFEB-STM.

2 Models

As previous research on the synchronization power of synchronization primitives [28], this paper assumes the linearizable shared memory model [6]. Due to NB-FEB combinability, as in [32] we assume that the processor-to-memory interconnection network is *nonovertaking* and that a reply message is sent back on the same path followed by the request message. The immediate nodes, on the communication path

Algorithm 1 *TFAS*(x : variable, v : value): Test-Flag-And-Set, a non-blocking variant of the original Store-if-Clear-and-Set primitive, which *always returns* the old value of x .

```

( $o, flag_o$ )  $\leftarrow$  ( $x, flag_x$ );
if  $flag_x = \text{false}$  then
  ( $x, flag_x$ )  $\leftarrow$  ( $v, \text{true}$ );
end if
return ( $o, flag_o$ );

```

Algorithm 2 *LOAD*(x : variable)

```

return ( $x, flag_x$ );

```

Algorithm 3 *SAC*(x : variable, v : value): Store-And-Clear

```

( $o, flag_o$ )  $\leftarrow$  ( $x, flag_x$ );
( $x, flag_x$ )  $\leftarrow$  ( $v, \text{false}$ );
return ( $o, flag_o$ );

```

Algorithm 4 *SAS*(x : variable, v : value): Store-And-Set

```

( $o, flag_o$ )  $\leftarrow$  ( $x, flag_x$ );
( $x, flag_x$ )  $\leftarrow$  ( $v, \text{true}$ );
return ( $o, flag_o$ );

```

from a processor to a global shared memory module (such as switches of a multistage interconnection network or higher memory modules of a multilevel memory hierarchy), can detect requests destined for the same destination and maintain the queues of requests. No memory coherent schemes are assumed.

3 NB-FEB Primitives

The set of NB-FEB primitives consists of four sub-primitives: *TFAS* (Algorithm 1), *Load* (Algorithm 2), *SAC* (Algorithm 3) and *SAS* (Algorithm 4). The last three primitives are similar to those of the original full/empty bit. Regarding conditional load primitives, a processor can check the flag value, $flag_x$, returned by the unconditional load primitive to determine if it was successful.

When the value of $flag_x$ returned is not needed, we just write $r \leftarrow \text{TFAS}(x, v)$ instead of $(r, flag_r) \leftarrow \text{TFAS}(x, v)$, where r is x 's old value. The same applies to *SAC* and *SAS*. For *Load*, we just write $r \leftarrow x$ instead of $r \leftarrow \text{LOAD}(x)$. In this paper, the flag value returned is needed only for combining NB-FEB primitives.

3.1 *TFAS*: A Universal Primitive

Lemma 1. (*Universality*) *The test-flag-and-set primitive (or TFAS for short) is universal.*

Proof. We will show that there is a wait-free¹ consensus algorithm, for arbitrary number of processes, that uses only the

¹An implementation is *wait-free* if it guarantees that any process can com-

Algorithm 5 TFAS_CONSENSUS(*proposal*: value)

Decision: shared variable. The shared variable is initialized to \perp with a clear flag (i.e. $flag_{Decision} = \text{false}$).

Output: a value agreed by all processes.
1T: $first \leftarrow \text{TFAS}(Decision, proposal)$;
2T: **if** $first = \perp$ **then**
3T: **return** $proposal$;
4T: **else**
5T: **return** $first$;
6T: **end if**

TFAS primitive and registers.

The wait-free consensus algorithm is shown in Algorithm 5. Processes share a variable called *Decision*, which is initialized to \perp with a *false* flag. Each process p proposes its value ($\neq \perp$) called *proposal* by calling TFAS_CONSENSUS(*proposal*).

The TFAS_CONSENSUS procedure is clearly wait-free since it contains no loops. We need to prove that i) the procedure returns the same value to all processes and ii) the value returned is the value proposed by some process. Indeed, the procedure will return the proposal of the first process executing TFAS on the *Decision* variable to all processes. Let p be a process calling the procedure.

- If p is the first process executing TFAS on the *Decision* variable, since the *Decision* variable is initialized to \perp with a *false* flag, p 's TFAS will successfully write p 's proposal to *Decision* and return \perp , the previous value of *Decision*. Since the value returned is \perp , the procedure returns p 's proposal (line 3T), the proposal of the first process executing TFAS.
- If p is not the first process executing TFAS on the *Decision* variable, p 's TFAS will fail to write p 's proposal to *Decision* since $flag_{Decision}$ has been set to *true* by the first TFAS on *Decision*. p 's TFAS will return the value, called *first*, written by the first TFAS. The *first* value is the proposal of the first process executing TFAS on the *Decision* variable. Since $first \neq \perp$ (due to the hypothesis that proposals are not \perp), the procedure will return *first* (line 5T).

□

3.2 Combinability

Lemma 2. (*Combinability*) NB-FEB primitives are combinable.

Proof. Table 1 summarizes the combining logic of NB-FEB primitives on a memory location x . The first column is the name of the first primitive request and the first row is the name

complete any operation on the implemented object in a finite number of steps, regardless of the execution speeds on the other processes [28, 34].

$(x, [v_1])$	The successive primitive with parameters $(x, [v_2])$			
	<i>Load</i>	<i>SAC</i>	<i>SAS</i>	<i>TFAS</i>
<i>Load</i>	<i>Load</i> (r, f_r) (r, f_r)	<i>SAC</i> (v_2) (r, f_r) (r, f_r)	<i>SAS</i> (v_2) (r, f_r) (r, f_r)	<i>TFAS</i> (v_2) (r, f_r) (r, f_r)
<i>SAC</i>	<i>SAC</i> (v_1) (r, f_r) $(v_1, 0)$	<i>SAC</i> (v_2) (r, f_r) $(v_1, 0)$	<i>SAS</i> (v_2) (r, f_r) $(v_1, 0)$	<i>SAS</i> (v_2) (r, f_r) $(v_1, 0)$
<i>SAS</i>	<i>SAS</i> (v_1) (r, f_r) $(v_1, 1)$	<i>SAC</i> (v_2) (r, f_r) $(v_1, 1)$	<i>SAS</i> (v_2) (r, f_r) $(v_1, 1)$	<i>SAS</i> (v_1) (r, f_r) $(v_1, 1)$
<i>TFAS</i>	<i>TFAS</i> (v_1) (r, f_r) Like 5th column	<i>SAC</i> (v_2) (r, f_r) Like 5th column	<i>SAS</i> (v_2) (r, f_r) Like 5th column	<i>TFAS</i> (v_1) (r, f_r) if $f_r=0$: $(v_1, 1)$ else: $(r, 1)$

Figure 1. The combining logic of NB-FEB primitives on a memory location x

of the successive primitive request. For instance, the cell [*SAS*, *TFAS*] is the combining logic of *SAS* and *TFAS* in which *SAS* is followed by *TFAS*. Let v_1, v_2, r and f_r be the value of the first primitive request, the value of the second primitive request, the value returned and the flag returned, respectively. In each cell, the first line is the combined request, the second is the reply to the first primitive request and the third (and fourth) is the reply to the successive primitive request. The values 0 and 1 of f_r in the reply represent *false* and *true*, respectively.

Consider the cell [*TFAS*, *TFAS*] as an example. The cell describes the case where request *TFAS*(x, v_1) is followed by request *TFAS*(x, v_2), at a switch of the processor-to-memory interconnection network. The two requests can be combined into only one request *TFAS*(x, v_1) (line 1), which will be forwarded further to the corresponding memory controller. When receiving a reply (r, f_r) to the combined request, the switch at which the requests were combined, creates separate replies to the two original requests. The reply to the first original request, *TFAS*(x, v_1), is (r, f_r) (line 2) as if the request was executed by the memory controller. The reply to the successive request, *TFAS*(x, v_2), depends on whether the combined request *TFAS*(x, v_1) has successfully updated the memory location x . If $f_r = 0$, *TFAS*(x, v_1) has successfully updated x with its value v_1 . Therefore, the reply to the successive request *TFAS*(x, v_2) is $(v_1, 1)$ as if the request was executed right after the first request *TFAS*(x, v_1). If $f_r = 1$, *TFAS*(x, v_1) has failed to update the x variable. Therefore, the reply to the successive request *TFAS*(x, v_2) is $(r, 1)$.

□

4 NBFEB-STM: Obstruction-free Multi-versioning STM

Like previous obstruction-free multi-versioning STM called LSA-STM [43], the new software transactional memory called NBFEB-STM, assumes that objects are only accessed and modified within transactions. NBFEB-STM assumes that there are no nested transactions, namely each thread executes only one transaction at a time. NBFEB-STM, like other obstruction-free STMs [30, 36, 43], is designed for garbage-collected programming languages (e.g. Java). A variable reclaimed by the garbage collector is assumed to have all bits 0 when it is reused. Note that there are non-blocking garbage collection algorithms that do not require synchronization primitives other than reads and writes while they still guarantee the non-blocking property for application-threads. Such a garbage collection algorithm is presented in Section 5.

Only two NB-FEB primitives, *TFAS* and *SAC*, are needed for implementing NBFEB-STM.

4.1 Challenges and Key Ideas

Unlike the STMs using *CAS* [30, 36, 43], NBFEB-STM using *TFAS* and *SAC* must handle the problem that *SAC*'s interference with concurrent *TFAS*es will violate the atomicity semantics expected on variable x . Overlapping *TFAS*₁ and *TFAS*₂ both may successfully write their new values to x if *SAC* interference occurs.

The key idea is not to use the transactional memory object *TMObj* [30, 36, 43] that needs to switch its pointer frequently to a new locator (when a transaction commits). Such a *TMObj* would need *SAC* in order to clear the pointer's flag, allowing the next transaction to switch the pointer. Instead, NBFEB-STM keeps a linked-list of locators for each object and integrates a write-once pointer *next* into each locator (cf. Figure 2). When opening an object O for write, a transaction T tries to append its locator to O 's locator-list by changing the *next* pointer of the head-locator of the list using *TFAS*. Due to the semantics of *TFAS*, only one of the concurrent transactions trying to append their locators succeeds. The other transactions must retry in order to find the new head and then append their locators to the new head. Using the locator-list, each *next* pointer is changed only once and thus its flag does not need to be cleared during the lifetime of the corresponding locator. This prevents a *SAC* from interleaving with concurrent *TFAS*es. The *next* pointer, together with its locator, will be reclaimed by the garbage collector when the lifetime of its locator is over. The garbage collector ensures that a locator will not be recycled until no thread/transaction has a reference to it.

Linking locators together creates another challenge on the space complexity of NBFEB-STM. Unlike the STMs using *CAS*, a delayed/halted transaction T in NBFEB-STM may prevent all locators appended after its locator in a locator-list from being reclaimed. As a result, T may make the system

run out of memory and thus prevent other transactions from making progress, violating the obstruction-freedom property. The key idea to solve the space challenge is to break the list of obsolete locators into pieces so that a delayed transaction T prevents from being reclaimed only the locator that T has a direct reference as in the STMs using *CAS*. The idea is based on the fact that only the head of O 's locator-list is needed for further accesses to the O object.

However, breaking the list of an obsolete object O also creates another challenge on finding the head of O 's locator-list. Obviously, we cannot use a head pointer as in non-blocking linked-lists since modifying such a pointer requires *CAS*. The key idea is to utilize the fact that there are no nested transactions and thus each thread has at most one *active* locator² in each locator list. Therefore, by recording the latest locator of each thread appended to O 's locator-list, a transaction can find the head of O 's locator list. The solution is elaborated further in Section 4.2 and Section 4.3.

Based on the key ideas, we come up with the data structure for a transactional memory object that is illustrated in Figure 2 and presented in Algorithm 6.

The transactional memory object in NBFEB-STM is an array of N pairs (pointer, timestamp), where N is the number of concurrent threads/transactions as shown in Figure 2. Item *TMObj*[i] is modified only by thread t_i and can be read by all threads. Pointer *TMObj*[i].*loc* points to the locator called *Loc* _{i} corresponding to the latest transaction committed/aborted by thread t_i . Timestamp *TMObj*[i].*ts* is the commit timestamp of the object referenced by *Loc* _{i} .*old*. After successfully appending its locator *Loc* _{i} to the list by executing *TFAS*(*head.next*, *Loc* _{i}), t_i will update its own item *TMObj*[i] with its new locator *Loc* _{i} . The *TMObj* array is used to find the head of the list of locators *Loc*₁, \dots , *Loc* _{N} .

For each locator *Loc* _{i} , in addition to fields *Tx*, *old* and *new* that reference the corresponding transaction object, the old data object and the new data object, respectively, as in DSTM[30], there are two other fields *cts* and *next*. The *cts* field records the commit timestamp of the object referenced by *old*. The *next* field is the pointer to the next locator in the locator list. The *next* pointer is modified by NB-FEB primitives. In Figure 2, values {0, 1} in the *next* pointer denote the values {*false*, *true*} of its flag, respectively. The *next* pointer of the head of the locator list, *Loc*₃.*next*, has its flag clear (i.e. 0), and the *next* pointers of previous locators (e.g. *Loc*₁.*next*, *Loc*₂.*next*) have their flags set (i.e. 1) since their *next* pointers were changed. The *next* pointer of a new locator (e.g. *Loc*₄.*next*) is initialized to (\perp , 0). Due to the garbage collector semantics, all locators *Loc* _{j} reachable from the *TMObj* *shared* object by following their *Loc* _{j} .*next* pointers, will not be reclaimed.

For each transaction object *Tx* _{i} , in addition to fields *status*, *readSet* and *writeSet* corresponding to the status, the set of objects opened for read, and the set of objects opened for

²An *active* locator is a locator that is still in use, opposite to an *obsolete* locator.

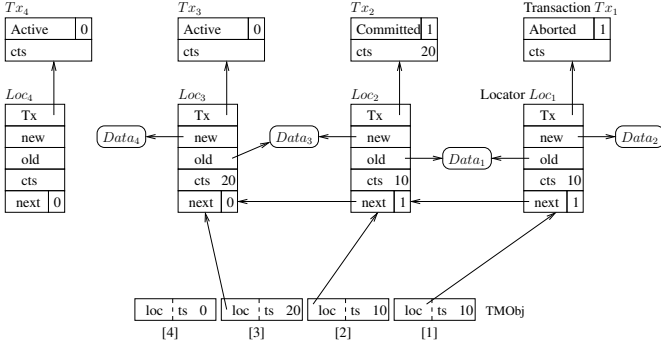


Figure 2. The data structure of a transactional memory object $TMOBJ$ in NBFEB-STM with four threads.

write, respectively, there is a field cts recording Tx_i 's commit timestamp (if Tx_i committed) as in LSA-STM [43].

4.2 Algorithm

A thread t_i starts a transaction T by calling the $STARTSTM(T)$ procedure (Algorithm 6). The procedure sets $T.status$ to *Active* and clears its flag using SAC (cf. Algorithm 3). The procedure then initializes the lazy snapshot algorithm (LSA) [43] by calling LSA_START . NBFEB-STM utilizes LSA to preclude inconsistent views by *live* transactions, an essential aspect of transactional memory semantics [25]. The LSA has been shown to be an efficient mechanism to construct consistent snapshots for transactions [43]. Moreover, the LSA can utilize up to $(N + 1)$ versions of an transactional memory object $TMOBJ$ recorded in N locators of $TMOBJ$'s locator list. Note that the global counter CT in LSA can be implemented by the *fetch-and-increment* primitive [22], a combinable (and thus scalable) primitive [32]. Except for the global counter CT , the LSA in NBFEB-STM does not need any strong synchronization primitives other than $TFAS$. The $ABORT(T)$ operation in LSA, which is used to abort a transaction T , is replaced by $TFAS(T.status, Aborted)$. Note that the *status* field is the only field of a transaction object T that can be modified by other transactions.

When a transaction T opens an object O for read, it invokes the $OPENR$ procedure (Algorithm 7). The procedure simply calls the LSA_OPEN procedure of LSA [43] in the *Read* mode to get the version of O that maintains a consistent snapshot with the versions of other objects being accessed by T . If no such a version of O exists, LSA_OPEN will abort T and consequently $OPENR$ will return \perp (line 3R). That means there is a conflicting transaction that makes T unable to maintain a consistent view of all the object being accessed by T . Otherwise, $OPENR$ returns the version of O that is selected by LSA. This version is guaranteed by LSA to belong to a consistent view of all the objects being accessed by T . Up to $(N + 1)$ versions are available for each object O in NBFEB-

Algorithm 6 $STARTSTM(T: \text{transaction})$

$TMOBJ$: array[N] of $\{ptr, ts\}$. Pointer $TMOBJ[i].ptr$ points to the locator called Loc_i corresponding to the latest transaction committed/aborted by thread t_i . Timestamp $TMOBJ[i].ts$ is the commit timestamp of the object referenced by $Loc_i.old$. N is the number of concurrent threads/transactions. $TMOBJ[i]$ is written only by thread t_i .

Locator: **record** tx, new, old : pointer; cts : timestamp; **end**. The cts timestamp is the commit timestamp of the old version.

Transaction: **record** $status$: $\{Active, Committed, Aborted\}$; cts : timestamp; **end**. NBFEB-STM also keeps read/write sets as in LSA-STM, but the sets are omitted from the pseudocode since managing the sets in NBFEB-STM is similar to LSA-STM.

1S: $SAC(T.status, Active)$; // Store-and-clear
 2S: $LSA_START(T)$ // Lazy snapshot algorithm

Algorithm 7 $OPENR(T: \text{Transaction}; O_i: TMOBJ)$: Open a transactional object for read

Output: reference to a *data* object if succeeds, or \perp .

1R: $LSA_OPEN(T, O_i, "Read")$; // LSA's OPEN procedure
 2R: **if** $T.status = Aborted$ **then**
 3R: **return** \perp ;
 4R: **else**
 5R: **return** the version chosen by LSA_OPEN ;
 6R: **end if**

STM (cf. Lemma 8). Since NBFEB-STM utilizes LSA, read-accesses to an object O are invisible to other transactions and thus do not change O 's locator list.

When a transaction T opens an object O for write, it invokes the $OPENW$ procedure (cf. Algorithm 8). The task of the procedure is to append to the head of O 's locator list a new locator L whose Tx and old fields reference to T and O 's latest version, respectively. In order to find O 's latest version, the procedure invokes $FINDHEAD$ (cf. Algorithm 9) to find the current head of O 's locator list (line 3W). When the head called H is found, the procedure determines O 's latest version based on the status of the corresponding transaction $H.Tx$ as in $DSTM$ [30]. If the $H.Tx$ transaction committed, O 's latest version is $H.new$ with commit timestamp $H.Tx.cts$ (lines 5W-7W). A copy of O 's latest version is created and referenced by $L.new$ (line 8W) (cf. locators Loc_2 and Loc_3 in Figure 2 as H and L , respectively, for an illustration). If the $H.Tx$ transaction aborted, O 's latest version is $H.old$ with commit timestamp $H.cts$ (lines 10W-12W) (cf. locators Loc_1 and Loc_2 in Figure 2 as H and L , respectively, for an illustration). If the $H.Tx$ transaction is active, $OPENW$ consults the contention manager [24, 50] (line 16W) to solve the conflict between the T and $H.Tx$ transactions. If T must

abort, OPENW tries to change $T.status$ to *Aborted* using $TFAS$ (line 18W) and returns \perp . Note that other transactions change $T.status$ only to *Aborted*, and thus if $TFAS$ at line 18W fails, $T.status$ has been changed to *Aborted* by another transaction. If $H.Tx$ must abort, OPENW changes $H.Tx.status$ to *Aborted* using $TFAS$ (line 21W) and checks $H.Tx.status$ again.

The latest version of O is then checked to ensure that it, together with the versions of other objects being accessed by T , belongs to a consistent view using LSA_OPEN with "Write" mode (line 28W). If it does, OPENW tries to append the new locator L to O 's locator list by changing the $H.next$ pointer to L (line 32W). Note that the $H.next$ pointer was initialized to \perp with a clear flag, before H was successfully appended to O 's locator list (line 27W). If OPENW does not succeed, another locator has been appended as a new head and thus OPENW must retry to find the new head (line 33W). Otherwise, it successfully appends the new locator L as the new head of O 's locator list. OPENW, which is being executed by a thread t_i , then makes $O[i].ptr$ reference to L and records $L.cts$ in $O[i].ts$ (line 36W). This removes O 's reference to the previous locator $oldLoc$ appended by t_i , allowing $oldLoc$ to be reclaimed by the garbage collector. Since $oldLoc$ now becomes an obsolete locator, its $next$ pointer is reset (line 37W) to break possible chains of obsolete locators reachable by a delayed/halted thread, helping $oldLoc$'s descendant locators in the chains be reclaimed. For each item j in the O array such that $O[j].ts < O[i].ts$, the $O[j].ptr$ locator now becomes obsolete in a sense that it no longer keeps O 's latest version although it is still referenced by $O[j]$ (since only thread t_j can modify $O[j]$). In order to break the chains of obsolete locators, OPENW resets the $next$ pointer of the $O[j].ptr$ locator so that $O[j].ptr$'s descendant locators can be reclaimed by the garbage collector (lines 38W-39W). This chain-breaking mechanism makes the space complexity of an object updated by N concurrent transactions/threads in NBFEB-STM be $\Theta(N)$, the optimal (cf. Theorem 1).

In order to find the head of O 's locator list as in OPENW, a transaction invokes the $FINDHEAD(O)$ procedure (cf. Algorithm 9). The procedure atomically reads O into a local array $start$ (line 2F). Such a multi-word read operation is supported by emerging multicore architectures like CUDA [39] and Cell BE [23]. In the contemporary chips of these architectures, a read operation can atomically read 128 bytes. In general, such a multi-word read operation can be implemented as an atomic snapshot using only single-word read and single-word write primitives [1]. $FINDHEAD$ finds the item $start_{latest}$ with the highest timestamp in $start$ and searches for the head from locator $start_{latest}.ptr$ by following the $next$ pointers until it finds a locator H whose $next$ pointer is \perp (lines 3F-6F). Since some locators may become obsolete and their $next$ pointers were reset to \perp by concurrent transactions (lines 37W and 39W in Algorithm 8), $FINDHEAD$ needs to check H 's commit timestamp against the highest timestamp of O at a moment after H is found (lines 8F-10F). If H 's commit timestamp is

Algorithm 8 OPENW(T : Transaction; O : TMObj): Open a transactional memory object for write by a thread p_i

Output: reference to a *data* object if succeeds, or \perp .

```

1W:  $newLoc \leftarrow$  new Locator;
2W: while true do
3W:    $head \leftarrow$  FINDHEAD( $O$ ); // Find the head of  $O$ 's list.
4W:   for  $i = 0$  to 1 do
5W:     if  $head.tx.status = Committed$  then
6W:        $newLoc.old \leftarrow head.new$ ;
7W:        $newLoc.cts \leftarrow head.tx.cts$ ;
8W:        $newLoc.new \leftarrow$  COPY( $head.new$ ); // Create a
           duplicate
9W:       break;
10W:     else if  $head.tx.status = Aborted$  then
11W:        $newLoc.old \leftarrow head.old$ ;
12W:        $newLoc.cts \leftarrow head.cts$ ;
13W:        $newLoc.new \leftarrow$  COPY( $head.old$ );
14W:       break;
15W:     else
16W:        $myProgression \leftarrow$  CM( $O_i, "Write"$ ) //
            $head.tx$  is active  $\Rightarrow$  Consult the contention
           manager
17W:       if  $myProgression = false$  then
18W:         TFAS( $T.status, Aborted$ ); // If fails, an-
           other has executed this  $TFAS$ .
19W:         return  $\perp$ ;
20W:       else
21W:         TFAS( $head.tx.status, Aborted$ );
22W:         continue; // Transaction  $head.tx$  has com-
           mitted/aborted  $\Rightarrow$  Check  $head.tx.status$  one
           more time
23W:       end if
24W:     end if
25W:   end for
26W:    $newLoc.tx \leftarrow T$ ;
27W:   SAC( $newLoc.next, \perp$ ); // Store-and-clear
28W:   LSA_OPEN( $T, O, "Write"$ ); // LSA's OPEN proce-
           dure.
29W:   if  $T.status = Aborted$  then
30W:     return  $\perp$ ; // Performance (not correctness): Don't
           add  $newLoc$  to  $O$  if  $T$  has aborted due to, for in-
           stance, LSA_OPEN.
31W:   end if
32W:   if TFAS( $head.next, newLoc$ )  $\neq \perp$  then
33W:     continue; // Another locator has been appended  $\Rightarrow$ 
           Find the head again
34W:   else
35W:      $oldLoc = O[i]$ ;
36W:      $O[i] \leftarrow (newLoc, newLoc.cts)$ ; // Atomic assign-
           ment;  $p_i$ 's old locator is unlinked from  $O$ .
37W:     SAC( $oldLoc.next, \perp$ ); //  $oldLoc$  may be in the
           chain of a sleeping thread  $\Rightarrow$  Stop the chain here
38W:     for each item  $L_j$  in  $O$  such that  $L_j.ts < O[i].ts$ 
           do
39W:       SAC( $L_j.ptr.next, \perp$ ) // Reset the  $next$  pointer
           of the obsolete locator
40W:     end for
41W:     return  $newLoc.new$ ;
42W:   end if
43W: end while

```

Algorithm 9 FINDHEAD(O : TMOBJ): Find the head of the locator list

Output: reference to the head of the locator list

```

1F: repeat
2F:    $start \leftarrow O$ ; // Read  $O$  to a local array atomically.
3F:   Let  $start_{latest}$  is the item with highest timestamp;
4F:    $tmp \leftarrow start_{latest}.ptr$ ; // Find a locator whose  $next$ 
      pointer is  $\perp$ 
5F:   while  $tmp.next \neq \perp$  do
6F:      $tmp \leftarrow tmp.next$ ;
7F:   end while
8F:    $start' \leftarrow O$ ; // Check if  $tmp$  is the head.
9F:   Let  $start'_{latest}$  is the item with highest timestamp;
10F:  until  $tmp.cts \geq start'_{latest}.ts$ ;
11F:  return  $tmp$ ;

```

Algorithm 10 COMMITW(T : Transaction): Try to commit an update transaction T by thread p_i

```

1C:  $CT_T \leftarrow LSA\_COMMIT(T)$ ; // Check consistent snapshot.  $CT_T$  is  $T$ 's unique commit timestamp from LSA.

2C:  $T.cts \leftarrow CT_T$ ; // Commit timestamp of  $T$  if  $T$  manages to commit.

3C:  $TFAS(T.status, Committed)$ ;

```

greater than or equal to the highest timestamp of O , H is the head of O 's locator list (cf. Lemma 4). Otherwise, H is an obsolete locator and FINDHEAD must retry (line 10F). The FINDHEAD procedure is lock-free, namely it will certainly return the head of O 's locator list after at most N iterations unless a concurrent thread has completed a transaction and subsequently has started a new one, where N is the number of concurrent (updating) threads (cf. Lemma 5). Note that as soon as a thread obtains $head$ from FINDHEAD (line 3W of OPENW, Algorithm 8), the locator referenced by $head$ will not be reclaimed by the garbage collector until the thread returns from the OPENW procedure.

When committing, read-only transactions in NBFEB-STM do nothing and always succeed in their commit phase as in LSA-STM [43]. They can abort only when trying to open an object for read (cf. Algorithm 7). Other transactions T , which have opened at least one object for write, invoke the COMMITW procedure (Algorithm 10). The procedure calls the LSA_COMMIT procedure to ensure that T still maintains a consistent view of objects being accessed by T (line 1C). T 's commit timestamp is updated with the timestamp returned from LSA_COMMIT (line 2C). Finally, COMMITW tries to change $T.status$ to *Committed* (line 3C). $T.status$ will be changed to *Committed* at this step if it has not been changed to *Aborted* due to the semantics of *TFAS*.

4.3 Analysis

In this section, we prove that NBFEB-STM fulfills the three essential aspects of transactional memory semantics [25]:

Instantaneous commit : Committed transactions must appear as if they executed instantaneously at some unique point in time, and aborted transactions, as if they did not execute at all.

Preserving real-time order : If a transaction T_i commits before a transaction T_j starts, then T_i must appear as if it executed before T_j . Particularly, if a transaction T_1 modifies an object O and commits, and then another transaction T_2 starts and reads O , then T_2 must read the value written by T_1 and not an older value.

Precluding inconsistent views : The state (of shared objects) accessed by *live* transactions must be consistent.

First, we prove some key properties of NBFEB-STM.

Lemma 3. A locator L_i with timestamp cts_i does not have any links/references to another locator L_j with a lower timestamp $cts_j < cts_i$.

Proof. There is only the *next* pointer to link between locators. The *next* pointer of locator L_i points to a locator L_j only if $L_j.cts$ is not less than $L_i.cts$ (lines 7W and 12W, Algorithm 8). Note that for each locator L_i , the commit timestamp $L_i.tx.cts$ of its corresponding transaction $L_i.tx$ (if $L_i.tx$ committed) is the commit timestamp of L 's new data and thus it is always greater than the commit timestamp $L_i.cts$ of L_i 's old data. \square

Lemma 4. The locator returned by FINDHEAD(O) (Algorithm 9) is the head H of O 's locator list at the time-point FINDHEAD found $H.next = \perp$ (line 5F).

Proof. Let L be the locator returned by FINDHEAD. Since the *next* pointer of a new locator is initialized to \perp (line 27W, Algorithm 8) before the locator is appended into the list by *TFAS* (line 32W), FINDHEAD will find a locator L whose *next* pointer is \perp at a time-point tp (line 5F). The L locator is either the head at that time or a reset locator (due to lines 37W and 39W, Algorithm 8).

If L is a reset locator, $start'_{latest}.cts > L.cts$ holds (line 10F) since a locator is reset (e.g. *oldLoc* at line 37W or L_j at line 39W) only after a locator with a higher timestamp (e.g. *newLoc*) has been written into the O array (line 36W). Since FINDHEAD atomically reads the O array after it found $L.next = \perp$, it will observe the higher timestamp. This makes FINDHEAD retry and discard L , a contradiction to the hypothesis that L is returned by FINDHEAD. Therefore, the L locator returned by FINDHEAD must be the head at the time-point FINDHEAD found $L.next = \perp$ (line 5F). \square

Since a thread must get a result from FINDHEAD (line 3W) before it can consult the contention manager (line 16W),

FINDHEAD must be lock-free (instead of being obstruction-free) in order to guarantee the obstruction-freedom for transactions.

Lemma 5. (Lock-freedom) *FINDHEAD(O) will certainly return the head of O 's locator list after at most N repeat-until iterations unless a concurrent thread has completed a transaction and subsequently has started a new one, where N is the number of concurrent threads updating O .*

Proof. From Lemma 4, any locator returned by FINDHEAD(O) is the head of O 's locator list. Therefore, we only need to prove that FINDHEAD(O) will certainly return a locator after at most N iterations unless a concurrent thread has completed a transaction and subsequently has started a new one.

We prove this by contradiction. Assume that FINDHEAD(O) executed by thread t_i , does not return after N iterations and no thread has completed its transaction since FINDHEAD started. Since each thread t_j updates its own item $O[j]$ only once when opening O for update (line 36W, , Algorithm 8), at most $(N - 1)$ items j of O , $j \neq i$, have been updated since FINDHEAD(O) started.

First we prove that FINDHEAD(O) will return in the iteration during which no item of O is updated between the first atomic read (line 2F) and the second atomic read of the O array (line 8F).

Indeed, since each transaction successfully appends its own locator to the head of O 's locator list only once when opening O for update (line 32W), at most $(N - 1)$ locators are appended to O 's locator list after the first scan. Therefore, FINDHEAD will certainly find a locator L such that $L.next \neq \perp$ (line 5F) in the current repeat-until iteration. Note that for each *next* pointer, only the first transaction executing TFAS on the pointer, manages to append its locator to the pointer.

Since (1) the *next* pointer of a locator L_i points to a locator L_j only if $L_j.cts \geq L_i.cts$ (cf. Lemma 3) and (2) FINDHEAD found L by following the *next* pointers starting from $start_{latest.ptr}$ (lines 3F-6F), we have $L.cts \geq start_{latest.ptr.cts}$. Note that $start_{latest.ptr.cts} = start_{latest.ts}$ (line 36W). Since no item of O is updated between the first scan (line 2F) and the second scan of the O array (line 8F), the items with highest timestamp of both scans are the same, i.e. $start_{latest} = start'_{latest}$. Therefore, $L.cts \geq start'_{latest.ts}$ holds (line 10F) and L is returned.

Since FINDHEAD executed by thread t_i does not return after N iterations due to hypothesis, it follows that at least N items have been updated since FINDHEAD started, a contradiction to the above argument that at most $(N - 1)$ items have been updated since FINDHEAD started. \square

Lemma 6. (Instantaneous commit) *TFAS-LSA guarantees that committed transactions appear as if they executed instantaneously and aborted transactions appear as if they did not execute at all.*

Proof. Similar to the DSTM [30] and LSA-STM [43], the NBFEB-STM uses the indirection technique that allows a trans-

action T_j to commit its modifications to all objects in its write-set instantaneously by switching its status from *Active* to *Committed*. Its committed status must no longer be changed. NBFEB-STM uses the TFAS primitive (Algorithm 1) to achieve the property (line 3C, Algorithm 10). Since the flag of the $T_j.status$ variable is *false* (or 0) when the transaction starts (line 1S, Algorithm 6), only the first TFAS primitive can change the variable. If T_j manages to change the $T_j.status$ variable to *Committed*, the variable is no longer able to be changed using TFAS until the transaction object T_j is reclaimed by the garbage collector. Note that even if thread t_j completed transaction T_j and has started another transaction T'_j , the transaction object T_j will not be reclaimed until all the locators keeping a reference to T_j are reclaimable.

Since active transactions T_j make all changes on their own copy $T_j.new$ of a shared object O before their status is changed from *Active* to either *Aborted* or *Committed*, aborted transactions do not affect the value of O . \square

The two other correctness criteria for transactional memory are precluding inconsistent views and preserving real-time order [25]. Since TFAS use the lazy snapshot algorithm LSA [43], the former will follow if we can prove that the LSA algorithm is integrated correctly into NBFEB-STM.

Lemma 7. *The versions kept in N locators $O[j].ptr$, $1 \leq j \leq N$, for each object O is enough for checking the validity of a transaction T using the LSA algorithm [43], from the correctness point of view.*

Proof. The LSA algorithm requires only the commit timestamp (i.e. $\lfloor O^{CT} \rfloor^3$) of the most recent version (i.e. O^{CT} ⁴) of each object O at a timestamp CT when it checks the validity of a transaction T . The older versions of O are not required for correctness - they only increase the chance that a suitable object version is available.

We will prove that by atomically reading the O object/array at the timestamp CT to a local variable V as at line 2F in Algorithm 9, LSA will find the commit timestamp $\lfloor O^{CT} \rfloor$.

A new version of O is created and becomes accessible by all transactions when a transaction T_j commits its modification $L_j.new$ (stored in locator L_j) to O by changing its status from *Active* to *Committed* (line 3C, Algorithm 10). Since every transaction T_j writes its locator L_j to $O[j].ptr$ when opening O for update (line 36W, Algorithm 8) (i.e. before committing), at least one of the locators $O[j].ptr$, $1 \leq j \leq N$, must contain the most recent version of O at the timestamp CT when O is read to V .

Since a transaction T_j updates $O[j]$ with its new locator L_j only after successfully appending L_j to the head of O 's locator list, at most one of the locators $O[j].ptr$, $1 \leq j \leq N$, is the head of the list at the timestamp CT when the snapshot V of O is taken. Other locators $V[j].ptr$ that are not the head, have their transactions committed/aborted before CT . Note

³Term $\lfloor O^t \rfloor$ denotes the time of most recent update of object O performed no later than time t [43].

⁴Term O^t denotes the content/version of object O at time t [43].

that as soon as the transaction of a locator committed/aborted, the locator's versions together with their commit timestamp is no longer changed. If transaction $V[i].ptr.tx$ committed, the version kept in locator $V[j].ptr$ is $V[j].ptr.new$ with commit timestamp $V[j].ptr.tx.cts$, the commit timestamp of the transaction. If transaction $V[j].ptr.tx$ has been aborted or is active, the version is $V[j].ptr.old$ with commit timestamp $V[j].ptr.cts$. The only possible version with commit timestamp higher than CT is $V[h].ptr.new$ where $V[h].ptr$ was the head at the timestamp CT when V was taken and then transaction $V[h].ptr.tx$ committed. In this case, $V[h].ptr.old$ is the most recent version at CT and its commit timestamp is $V[h].ptr.cts$.

Therefore, by checking the commit timestamps of the versions kept in each locator $V[j].ptr$, $1 \leq j \leq N$, against CT , LSA will find the commit timestamp $\lfloor O^{CT} \rfloor$ of the most recent update of object O performed no later than CT . \square

Lemma 8. *The number of versions available for each object in NBFEB-STM is up to $(N + 1)$, where N is the number of threads.*

Proof. For each object O , each thread t_j keeps a version of O that has been accessed most recently by t_j , in locator $O[j].ptr$ (or L_j for short). If t_j 's latest transaction T_j committed $\forall j \in [1, N]$, the $L_j.old$ is an old version of O with validity range $[L_j.cts, L_j.tx.cts]$ ⁵. Therefore, if every thread has its latest transaction committed, each object O updated by N threads will have N old versions with validity ranges, additional to its latest version. \square

Lemma 9. (Consistent view) *NBFEB-STM precludes inconsistent views of shared objects from live transactions.*

Proof. Since the LSA lazy snapshot algorithm is correctly integrated into NBFEB-STM (Lemma 7), the lemma follows. \square

Definition 1. *The value of a locator L is either $L.new$ if $L.tx.status = Committed$, or $L.old$ otherwise.*

Lemma 10. *In each O 's locator list, the old value $L'.old$ of a locator L' is not older than the value of its previous locator⁶ L .*

Proof. Let L'' be the locator pointed by $L.next$. Since $L.tx.status$ must be either *Committed* or *Aborted* (but not *Active*) before L'' is appended to $L.next$ (lines 5W-24W, Algorithm 8), $L''.old$ is L 's value, which is either $L.new$ if $L.tx.status = Committed$ (line 6W) or $L.old$ if $L.tx.status = Aborted$ (line 11W). That means $L''.old$ is not older than L 's value. Arguing inductively for all locators on the directed path from L to L' , the lemma follows. \square

⁵The *validity range* of a version v_i of an object O is the interval from the commit time of v_i to the commit time of the next version v_{i+1} of O [43].

⁶A locator L is a *previous* locator of a locator L' if starting from L we can reach L' by following *next* pointers.

Lemma 11. (Real-time order preservation) *NBFEB-STM preserves the real-time order of transactions.*

Proof. We need to prove that if a transaction T_1 modifies an object O and commits and then another transaction T_2 starts and reads O , T_2 must read the value written by T_1 and not an older value [25]. Namely, T_1 is the most recent transaction committing its modification to O before T_2 reads O .

First we prove that T_2 reads the value v_1 written by T_1 if T_2 opens O for read (cf. OPENR, Algorithm 7). In the proof of Lemma 7, we have proven that the value of O read at a timestamp CT by LSA is the most recent value of O at that timestamp. Since T_1 is the most recent transaction committing its modification to O before T_2 reads O , v_1 is in the set of available versions of O read by LSA_OPEN (line 1R). Since T_1 commits before T_2 starts and reads O , the commit timestamp of v_1 is less than the upper bound of any validity range R_{T_2} ⁷ chosen by the LSA_OPEN (i.e. $\lfloor O^{CT} \rfloor \leq T_{max}$ in terminology used by LSA [43].) Therefore, the LSA_OPEN in OPENR will return v_1 , which is subsequently returned by OPENR (line 5R).

We now prove that T_2 reads the value v_1 written by T_1 if T_2 opens O for read (cf. OPENW, Algorithm 8). Particularly, we prove that the *old* value of T 's new locator (lines 6W and 11W) is v_1 .

Let p_1 and p_2 be the threads executing T_1 and T_2 , respectively, L_1 be the locator containing T_1 's modification (in $L_1.new$) that is committed to O and v_2 be the value of O read by T_2 . The v_2 value is the value of the head H of O 's locator list returned from FINDHEAD executed by T_2 , which is either $H.new$ if $H.ts.status = Committed$ or $H.old$ otherwise (line 6W or 11W).

Since T_1 committed before T_2 started, H is the head of O 's locator list that includes L_1 (cf. Lemma 4). Note that since T_1 is the latest transaction committing its modification to O , all locators L' that have ever been reachable from L_1 via *next* pointers, have the most recent timestamp/value (cf. Lemma 10) and thus will not be reset (lines 38W-39W, Algorithm 8). Since there is a directed path from L_1 to H via *next* pointers, it follows from Lemma 10 that the value of H is not older than that of L_1 .

On other hand, since T_1 is the latest transaction committing its modification to O before T_2 reads O , there is no value of O that is newer than that of L_1 . Therefore, the value of H is the value of L_1 . That means T_2 reads the v_1 value written by T_1 .

Finally, we need to prove that LSA_OPEN at line 28W accepts v_1 . Indeed, since v_1 is the most recent update of O and T_1 commits before T_2 starts, the commit timestamp of v_1 is less than the upper bound of any validity range R_{T_2} chosen by the LSA_OPEN (i.e. $\lfloor O^{CT} \rfloor \leq T_{max}$). Therefore, the LSA_OPEN at line 28W accepts v_1 . \square

⁷The *validity range* R_T of a transaction T is the time range during which each of the objects accessed by T is valid [43].

Lemma 12. *For each object O , there are at most $4N$ locators that cannot be reclaimed by the garbage collector at any time-point, where N is the number of update threads.*

Proof. Let L_i be a locator created by a thread p_i . A locator L_i cannot be reclaimed by the garbage collector if it is reachable by a thread. In NBFEB-STM, a locator L_i is reachable if it is i) p_i 's new locator $newLoc$, ii) p_i 's shared locator, which is referenced directly by $O[i].ptr$, and iii) p_i 's old locators $oldLoc$ that is reachable by other threads. p_i 's shared locator will become one of p_i 's old locators if $O[i].ptr$ is updated with p_i 's new locator (line 36W, Algorithm 8). At that moment, p_i 's new locator becomes p_i 's shared locator. If there is no thread keeping a direct/indirect reference to p_i 's old locators, these locators are ready to be reclaimed (i.e. unreachable) when p_i returns from the OPENW procedure.

Let C_i^p and C_i^o be the chains of locators (linked by their *next* pointers) that cannot be reclaimed due to thread p_i and $O[i]$, respectively. The C_i^p chain starts at the locator that is referenced directly by p_i (not directly by O) and ends at either the locator whose *next* pointer is \perp or the locator whose next locator is referenced directly by another thread or O . The C_i^o chain starts at the locator that is referenced directly by $O[i]$ and ends at either the locator whose *next* pointer is \perp or the locator that is referenced directly by another thread or O . Note that there are no two locators whose *next* pointers point to the same locator L_j since p_j successfully appends L_j into the head of the locator list only once (line 32W, Algorithm 8).

At any time, each thread p_i has at most one C_i^p and one C_i^o . The C_i^p starts either with p_i 's new locator (before assignment $O[i] \leftarrow newLoc$ at line 36W, Algorithm 8) or with p_i 's old locator (after this assignment). Since p_i has a unique item in the O array, it has at most one C_i^o . Therefore, there are at most $2N$ chains.

We will prove that if p_i has three locators participating in chains (of arbitrary threads), at least one of the three locators must be the end-locator of a chain. Indeed, during the execution of the OPENW procedure (Algorithm 8), p_i creates only one new locator (line 1W) in addition to its locator $O[i].ptr$, if any. If p_i has three locators that are participating in chains, at least one of them is p_i 's old locator L^o resulting from one of p_i 's previous executions E of OPENW. Since p_i sets the *next* pointer of its old locator $oldLoc$ to \perp before returning from E (line 37W), L^o 's *next* pointer is \perp . That means L^o is the end-locator of a chain.

It then follows that each thread has at most two *non-end* locators participating in all the chains. The number of non-end locators in all the chains is at most $2N$. Since there are at most $2N$ chains, there are at most $2N$ end-locators. Therefore, the total number of locators in all the chains is $4N$. \square

Theorem 1. (Space complexity) *The space complexity of an object updated by N threads in NBFEB-STM is $\Theta(N)$, the optimal.*

Proof. Since each object O in NBFEB-STM is an array of N items (cf. Algorithm 6), the space complexity of an object is

$\Omega(N)$.

From Lemma 12, for each object O there are at most $4N$ locators that cannot be reclaimed by the garbage collector at any point in time. Since each locator L references to at most one transaction object $L.tx$ (cf. Figure 2), the space complexity of an object is $O(N)$.

Due to the *instantaneous commit* requirement of transactional memory semantics [25], when opening an object for update, each thread/transaction in any STM system must create a copy of the original object. Therefore, the space complexity of an object updated by N threads is $O(N)$ for all STM systems. It follows that the space complexity $\Theta(N)$ of an object updated by N threads in NBFEB-STM is optimal. \square

Definition 2. *Contention level $CL_{l,t}$ of a memory location l at a timestamp t is the number of requests that need to be executed sequentially on the location by a memory controller (i.e. the number of requests for l buffered at time t).*

Definition 3. *Contention level of a transaction T that starts at timestamp s_T and ends (i.e. commits or aborts) at timestamp e_T is $\max_{s_T \leq t \leq e_T} CL_{l,t}$ for all memory locations l accessed by T*

Lemma 13. (Contention reduction) *Transactions using NBFEB-STM have lower contention levels than those using CAS-based STMs do.*

Proof. (Sketch) Since CAS is not combinable [32, 10], M conflicting CAS primitives on the same synchronization variable, like $TMObj$ pointer or a transaction's *status* variable in CAS-based STMs [30, 36, 43], issue M remote-memory requests to the corresponding memory controller. Since TFAS is combinable, the remote-memory requests from M conflicting TFAS primitives to the same variable, like the *next* pointer or a transaction's *status* variable in NBFEB-STM, can be combined into only one request to the corresponding memory controller. Therefore, the combinable primitive significantly reduces the number of requests for each memory location buffered at the memory controller. \square

5 Garbage Collectors

In this section, we present a non-blocking garbage collection algorithm called NB-GC that can be used in the context of NBFEB-STM. The NB-GC algorithm does not require synchronization primitives other than reads and writes while it still guarantees the obstruction-freedom property for *application threads* (or mutators in the memory management terminology). The obstruction-freedom here means that a halted application-thread cannot prevent other application-threads from making progress.

Like previous concurrent garbage collection algorithms for multiprocessors [4, 7, 8, 11, 13, 16, 18, 17, 19, 33, 35, 44, 46, 47, 26], the new NB-GC algorithm is a priority-based garbage

collection algorithm in which the collector thread is a privileged thread that may suspend and subsequently resume the mutator threads. The NB-GC algorithm is an improvement of the seminal on-the-fly garbage collector [16, 17, 18] using the sliding view technique [35] called SV-GC. Unlike the SV-GC algorithm, the NB-GC algorithm allows the collector to suspend a mutator at any point in the mutator’s code (even in the reference slot update and object allocation procedures). This prevents a mutator from blocking the collector and consequently from blocking other mutators.

In the concurrent garbage collection model, there are two kind of threads: application threads (e.g. the mutators) that perform user programs (error-prone codes), and privileged threads with higher priority (e.g. the collector) that perform system tasks (error-free codes). Whereas the application threads can be delayed/preempted arbitrarily, the system threads when running will not be preempted by the application threads. NB-GC guarantees obstruction-freedom for *application threads*, which usually perform users error-prone codes. Namely, a halted application-thread will not prevent other application-threads from making progress via blocking the garbage collector. The model, in some sense, covers the non-blocking garbage collection algorithms [29, 38] that, at the first look, seem not to require privileged threads. In fact, the non-blocking garbage collectors require strong synchronization primitives like *compare-and-swap* whose atomicity is guaranteed by hardware threads, a kind of privileged threads.

The SV-GC algorithm using the sliding view technique [35] does not need synchronization primitives other than reads and writes. However, it requires that the mutator be suspended only at a safe point, particularly it requires that the mutator not be stopped during the execution of a reference slot update nor new object allocation. If a mutator M is preempted during such an execution, the collector cannot progress since it cannot suspend the mutator M . This would prevent the other mutators from making progress due to lack of memory. Therefore, the SV-GC collector does not guarantee the obstruction-freedom for mutators and must rely heavily on the scheduler to avoid such a scenario.⁸

The basic idea of the sliding view technique in the SV-GC algorithm is as follows. At the beginning of a collection cycle k , the collector takes an asynchronous heap snapshot S_k of all (heap) reference slots s . By comparing snapshot S_{k-1} and S_k , the collector knows which objects have their reference counter changed during the interval between the two collections. For instance, if in the interval a reference slot s is sequentially assigned references to objects o_0, o_1, \dots, o_n , where (s, o_1) is recorded in S_{k-1} and (s, o_n) in S_k , the collector only needs to execute two reference count updates for o_0 and o_n : $RC(o_0) - -$ and $RC(o_n) + +$, instead of $2n$ reference count updates for o_0, o_n and $(n - 1)$ immediate objects $o_i, 1 \leq i \leq (n - 1)$: $RC(o_0) - -, RC(o_1) + +, RC(o_1) -$

⁸In order to reclaim unreachable cyclic structures of objects, the reference-counting collectors use either a backup tracing collector [7] infrequently or a cycle collector [40]. Both the efficient backup tracing collector [7] and cycle collector [40] use the sliding view technique.

Algorithm 11 GENERICCOLLECTOR: the main stages of a collection cycle using the sliding view technique

- 1: Raise the $Snoop_i$ flag of each mutator;
 - 2: Obtain a sliding view (concurrently with mutator’s computation);
 - 3: For each mutator M_i : 1) Suspend M_i ; 2) Turn the $Snoop_i$ flag off; 3) Mark as *local* objects O directly reachable from M_i ’s roots; 4) Resume M_i ;
 - 4: Update the reference counter $O.rc$ of each object O ;
 - 5: Reclaim objects O that are not marked *local* and $O.rc = 0$; For each descendent D of a reclaimed object, $D.rc - -$; D is checked for reclamation like O . This operation continues recursively until there are no objects that can be reclaimed.
-

$-, \dots, RC(o_n) + +$. The main stages of the *generic* sliding view algorithm [35] are shown in Algorithm 11. The algorithm is *generic* in the sense that it may use any mechanism for obtaining the sliding view. Instead of using an atomic snapshot algorithm [1] to obtain a consistent view of all heap reference slots, the algorithm uses a much simpler mechanism called *snooping* [16] to avoid wrong reference counts that result from an inconsistent view. For instance, if the only reference to an object O is moving from slot s_1 to slot s_2 when the view is taken, the view may miss the reference in both s_1 (reading after modification) and s_2 (reading before modification). To deal with the problem, the snooping mechanism marks as *local* any object that is assigned a new reference in the heap while the view is being read from the heap. The marked objects are left to be collected in the next collection cycle. The reader is referred to [35] for the complete SV-GC algorithm.

We found that the SV-GC algorithm [35] can be easily improved to provide obstruction-freedom for mutators using the *helping technique* [9]. Basically, if the collector suspends a mutator during its execution of a reference slot update or object allocation procedure, the collector helps the mutator by completing the procedure on behalf of the mutator and moving the mutator’s program counter (PC) to the end of the procedure before resuming the mutator. Note that in the concurrent garbage collection model there is only one collector that can suspend a given mutator and the collector suspends only one mutator at a time. The improved algorithm provides obstruction-freedom for mutators (or application-threads) by preventing mutators from blocking the collector and consequently from blocking other mutators. It is obstruction-free in the sense that progress is guaranteed for each active mutator regardless of the status of the other mutators.

Acknowledgments Phuong Ha’s and Otto Anshus’s work was supported by the Norwegian Research Council (grant numbers 159936/V30 and 155550/420). Philippas Tsigas’s work was supported by the Swedish Research Council (VR) (grant number 37252706).

References

- [1] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, and N. Shavit. Atomic snapshots of shared memory. *J. ACM*, 40(4):873–890, 1993.
- [2] A. Agarwal, J. Kubiawicz, D. Kranz, B.-H. Lim, D. Yeung, G. D’Souza, and M. Parkin. Sparcle: An evolutionary processor design for large-scale multiprocessors. *IEEE Micro*, 13(3):48–61, 1993.
- [3] R. Alverson, D. Callahan, D. Cummings, B. Koblenz, A. Porterfield, and B. Smith. The tera computer system. *SIGARCH Comput. Archit. News*, 18(3b):1–6, 1990.
- [4] A. W. Appel. Real-time concurrent collection on stock multiprocessors. *SIGPLAN Not.*, 39(4):205–216, 2004.
- [5] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick. The landscape of parallel computing research: A view from berkeley. *Technical Report No. UCB/EECS-2006-183, University of California, Berkeley*, 2006.
- [6] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. John Wiley and Sons, Inc., 2004.
- [7] H. Azatchi, Y. Levanoni, H. Paz, and E. Petrank. An on-the-fly mark and sweep garbage collector based on sliding views. In *Proc. of the ACM Conf. on Object-oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 269–281, 2003.
- [8] D. F. Bacon, C. R. Attanasio, H. B. Lee, V. T. Rajan, and S. Smith. Java without the coffee breaks: a non-intrusive multiprocessor garbage collector. In *Proc. of the ACM Conf. on Programming Language Design and Implementation (PLDI)*, pages 92–103, 2001.
- [9] G. Barnes. A method for implementing lock-free shared-data structures. *Proc. of the ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 261–270, 1993.
- [10] G. E. Blelloch, P. B. Gibbons, and S. H. Vardhan. Combinable memory-block transactions. In *Proc. of the ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 23–34, 2008.
- [11] H.-J. Boehm, A. J. Demers, and S. Shenker. Mostly parallel garbage collection. In *Proc. of the ACM Conf. on Programming Language Design and Implementation (PLDI)*, pages 157–164, 1991.
- [12] J. G. Castanos, L. Ceze, K. Strauss, and H. S. W. Jr. Evaluation of a multithreaded architecture for cellular computing. In *HPCA ’02: Proceedings of the 8th International Symposium on High-Performance Computer Architecture*, page 311, 2002.
- [13] P. Cheng and G. E. Blelloch. A parallel, real-time garbage collector. In *Proc. of the ACM Conf. on Programming Language Design and Implementation (PLDI)*, pages 125–136, 2001.
- [14] D. E. Culler, J. P. Singh, and A. Gupta. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann, 1998.
- [15] W. J. Dally, J. A. S. Fiske, J. S. Keen, R. A. Lethin, M. D. Noakes, P. R. Nuth, R. E. Davison, and G. A. Fyler. The message-driven processor: A multicomputer processing node with efficient mechanisms. *IEEE Micro*, 12(2):23–39, 1992.
- [16] E. W. Dijkstra, L. Lamport, A. J. Martin, C. S. Scholten, and E. F. M. Steffens. On-the-fly garbage collection: an exercise in cooperation. *Commun. ACM*, 21(11):966–975, 1978.
- [17] D. Doligez and G. Gonthier. Portable, unobtrusive garbage collection for multiprocessor systems. In *Proc. of the ACM Symp. on Principles of Programming Languages (POPL)*, pages 70–83, 1994.
- [18] D. Doligez and X. Leroy. A concurrent, generational garbage collector for a multithreaded implementation of ml. In *Proc. of the ACM Symp. on Principles of Programming Languages (POPL)*, pages 113–123, 1993.
- [19] T. Domani, E. K. Kolodner, and E. Petrank. A generational on-the-fly garbage collector for java. In *Proc. of the ACM Conf. on Programming Language Design and Implementation (PLDI)*, pages 274–284, 2000.
- [20] J. Feo, D. Harper, S. Kahan, and P. Konecny. Eldorado. In *CF ’05: Proceedings of the 2nd conference on Computing frontiers*, pages 28–34, 2005.
- [21] K. Fraser and T. Harris. Concurrent programming without locks. *ACM Trans. Comput. Syst.*, 25(2):5, 2007.
- [22] A. Gottlieb, R. Grishman, C. P. Kruskal, K. P. McAuliffe, L. Rudolph, and M. Snir. The nyu ultracomputer—designing a mimd, shared-memory parallel machine (extended abstract). *SIGARCH Comput. Archit. News*, 10(3):27–42, 1982.
- [23] M. Gschwind, H. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki. Synergistic processing in cell’s multicore architecture. *Micro, IEEE*, 26(2):10–24, 2006.
- [24] R. Guerraoui, M. Herlihy, and B. Pochon. Polymorphic contention management. In *Proc. of the Intl. Symp. on Distributed Computing (DISC)*, pages 303–323, 2005.
- [25] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *Proc. of the ACM Symp. on Principles and Practice of Parallel Programming (PPoPP)*, pages 175–184, 2008.

- [26] J. Guy L. Steele. Multiprocessing compactifying garbage collection. *Commun. ACM*, 18(9), 1975.
- [27] T. Harris and K. Fraser. Language support for lightweight transactions. In *Proc. of the ACM Conf. on Object-oriented Programing, Systems, Languages, and Applications (OOPSLA)*, pages 388–402, 2003.
- [28] M. Herlihy. Wait-free synchronization. *ACM Transaction on Programming and Systems*, 11(1):124–149, Jan. 1991.
- [29] M. Herlihy, V. Luchangco, P. Martin, and M. Moir. Non-blocking memory management support for dynamic-sized data structures. *ACM Trans. Comput. Syst.*, 23(2):146–196, 2005.
- [30] M. Herlihy, V. Luchangco, M. Moir, and I. William N. Scherer. Software transactional memory for dynamic-sized data structures. In *Proc. of Symp. on Principles of Distributed Computing (PODC)*, pages 92–101, 2003.
- [31] S. W. Keckler, W. J. Dally, D. Maskit, N. P. Carter, A. Chang, and W. S. Lee. Exploiting fine-grain thread level parallelism on the mit multi-alu processor. In *Proc. of the Intl. Symp. on Computer Architecture (ISCA)*, pages 306–317, 1998.
- [32] C. P. Kruskal, L. Rudolph, and M. Snir. Efficient synchronization of multiprocessors with shared memory. *ACM Trans. Program. Lang. Syst.*, 10(4):579–601, 1988.
- [33] L. Lamport. Garbage collection with multiple processes: an exercise in parallelism. In *Proc. of the Intl. Conf. on Parallel Processing*, pages 50–54, 1976.
- [34] L. Lamport. Concurrent reading and writing. *Commun. ACM*, 20(11):806–811, 1977.
- [35] Y. Levanoni and E. Petrank. An on-the-fly reference-counting garbage collector for java. *ACM Trans. Program. Lang. Syst.*, 28(1):1–69, 2006.
- [36] V. J. Marathe, W. N. S. Iii, and M. L. Scott. Adaptive software transactional memory. In *Proc. of the Intl. Symp. on Distributed Computing (DISC)*, pages 354–368, 2005.
- [37] M. Michael and M. Scott. Implementation of atomic primitives on distributed shared memory multiprocessors. In *High-Performance Computer Architecture, 1995. Proceedings., First IEEE Symposium on*, pages 222–231, 1995.
- [38] M. M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Trans. Parallel Distrib. Syst.*, 15(6):491–504, 2004.
- [39] NVIDIA. *NVIDIA CUDA Compute Unified Device Architecture, Programming Guide, version 1.1*. NVIDIA Corporation, 2007.
- [40] H. Paz, D. F. Bacon, E. K. Kolodner, E. Petrank, and V. T. Rajan. An efficient on-the-fly cycle collection. *ACM Trans. Program. Lang. Syst.*, 29(4):20, 2007.
- [41] G. F. Pfister, W. C. Brantley, D. A. George, S. L. Harvey, W. J. Kleinfelder, K. P. McAuliffe, E. S. Melton, V. A. Norton, and J. Weiss. The ibm research parallel processor prototype (rp3): Introduction and architecture. In *ICPP*, pages 764–771, 1985.
- [42] S. A. Plotkin. Sticky bits and universality of consensus. In *Proc. of Symp. on Principles of Distributed Computing (PODC)*, pages 159–175, 1989.
- [43] T. Riegel, P. Felber, and C. Fetzer. A lazy snapshot algorithm with eager validation. In *Proc. of the Intl. Symp. on Distributed Computing (DISC)*, pages 284–298, 2006.
- [44] J. Singer, G. Brown, I. Watson, and J. Cavazos. Intelligent selection of application-specific garbage collectors. In *Proc. of the Intl. Symp. on Memory management (ISMM)*, pages 91–102, 2007.
- [45] B. Smith. The architecture of hep. In *Parallel MIMD Computation: HEP Supercomputer and Its Applications, Scientific Computation Series*, page 41–55, 1985.
- [46] S. Soman, L. Daynès, and C. Krintz. Task-aware garbage collection in a multi-tasking virtual machine. In *ISMM '06: Proc. of the Intl. Symp. on Memory Management*, pages 64–73, 2006.
- [47] S. Soman and C. Krintz. Application-specific garbage collection. *J. Syst. Softw.*, 80(7):1037–1056, 2007.
- [48] S. Sridharan, A. Rodrigues, and P. Kogge. Evaluating synchronization techniques for light-weight multi-threaded/multicore architectures. In *Proc. of the ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 57–58, 2007.
- [49] C. Technology. Csx processor architecture whitepaper. 2006.
- [50] I. William N. Scherer and M. L. Scott. Advanced contention management for dynamic software transactional memory. In *Proc. of Symp. on Principles of Distributed Computing (PODC)*, pages 240–248, 2005.
- [51] W. Zhu, V. C. Sreedhar, Z. Hu, and G. R. Gao. Synchronization state buffer: supporting efficient fine-grain synchronization on many-core architectures. In *Proc. of the Intl. Symp. on Computer Architecture (ISCA)*, pages 35–45, 2007.