

Lock-free Parallel Garbage Collection by Mark&Sweep

Gao, H.¹, Groote, J.F.², Hesselink, W.H.¹

¹ Department of Mathematics and Computing Science, University of Groningen,
P.O. Box 800, 9700 AV Groningen, The Netherlands (Email: {hui,wim}@cs.rug.nl)

² Department of Mathematics and Computing Science, Eindhoven University of Technology, P.O. Box 513,
5600 MB Eindhoven, The Netherlands (Email: jfg@win.tue.nl)

Abstract

This paper presents a lock-free parallel algorithm for mark&sweep garbage collection (GC) in a realistic model using synchronization primitives *load-linked (LL)/store-conditional (SC)* or *compare-and-swap (CAS)* offered by machine architectures. *Mutators* and *collectors* can simultaneously operate on the data structure. In particular no strict alternation between usage and cleaning up is necessary contrary to what is common in most other garbage collection algorithms.

We first design and prove an algorithm with a coarse grain of atomicity and subsequently apply the reduction theorem developed in [17] to implement the higher-level atomic steps by means of the low-level primitives. Even so, the structure of our algorithm and its correctness properties, as well as the complexity of reasoning about them, makes neither automatic nor manual verification feasible. We have therefore chosen the higher-order interactive theorem prover PVS for mechanical support.

1 Introduction

On shared-memory multiprocessors, processes coordinate with each other via shared data structures. To ensure the consistency of these concurrent objects, processes need a mechanism for synchronizing their access. In such a system the programmer typically has to explicitly synchronize access to shared data by different processes to ensure correct behavior of the overall system, using synchronization primitives such as semaphores, monitors, guarded statements, mutex locks, etc. In fact, the operations of different processes on a shared data structure should appear to be serialized¹ so that the object state is kept coherent after each operation.

Due to blocking, the classical synchronization paradigms using locks can incur many problems such as convoying, priority inversion and deadlock. A *lock-free* (also called *non-blocking*) implementation of a shared object guarantees that within a finite number of steps always some process trying to perform an operation on the object will complete its task, independently of the activity and speed of other processes [21]. As lock-free synchronizations are built without locks, they are immune from the aforementioned problems. In addition, lock-free synchronizations can offer progress guarantees. A number of researchers [3, 5, 21, 22, 36, 38] have proposed techniques for designing lock-free implementations. Essential for such implementations are advanced machine instructions such as *load-linked (LL)/store-conditional (SC)*, or *compare-and-swap (CAS)*.

In this paper we propose a lock-free implementation of mark&sweep *garbage collection* (GC). Garbage *collectors* are employed to identify at run-time which objects are no longer referenced by the *mutators* (i.e. user programs that use and modify the objects). The heap space occupied by these objects is said to be *garbage* and must be re-cycled for subsequent new objects. The garbage collectors reclaim all garbage by adding them to a so called *free-list*, which keeps track of free memory. Some programming languages (e.g. C, C++) force or allow the programs to do their own GC, which means that programs are required to delete objects that they allocate in memory. However, this task is so difficult that non-trivial applications often exhibit incorrect behavior as the result of memory leaks or dangling pointers. To relieve programmers of virtually all memory-management problems, it is

¹if two operations execute simultaneously, the system guarantees the same result as if one of them is arbitrarily executed before the other.

preferable to offer GC that is automatically triggered during memory allocation when the amount of free memory falls below some threshold or after a certain number of allocations.

There are several basic strategies for GC: reference counting [10, 30, 35, 43], mark&sweep [1, 4, 11, 12, 13] and copying [23, 26, 27, 28, 44]. Reference counting algorithms can do their job incrementally (the entire heap need not be collected at once, resulting in shorter collection pauses), but impose overhead on the mutators and fail to reclaim circular garbage. Mark&sweep algorithms can reclaim circular structure, and don't place any burden on the mutators like reference counting algorithms do, but tend to leave the heap fragmented. Copying algorithms can reduce fragmentation, but add the cost of copying data from one space to another and require twice as much memory as a mark&sweep collector. Moreover, copying also requires that the programming language restrict address manipulation operations, which isn't true for C or C++. For a more detailed introduction to garbage collection and memory management the reader is referred to [31].

One often encounters GC algorithms (e.g. [6, 13, 14, 47]) that employ "stop-the-world" mechanisms, which suspend all normal running threads and then perform GC. Such an algorithm introduces a global synchronization point between all threads and tends to become a scaling bottleneck that limits program performance and processor utilization. In particular, a "stop-the-world" mechanism violates non-blockingness. This is unacceptable when the system must guarantee response time of interactive applications. Therefore, to achieve parallel speed-ups on shared-memory multiprocessors, lock-free algorithms are of interest [33, 48, 49].

There are several lock-free GC algorithms in the literature. The first one is due to Herlihy and Moss [23]. They present a lock-free copying GC algorithm, which uses copying for moving objects to avoid blocking synchronization. In their algorithm, the failure of a participating thread can indefinitely prevent the freeing of unbounded memory. In [26], Hesselink and Groote give a wait-free (wait-freedom is stronger than lock-freedom) GC algorithm using reference counting. However, this collector applies only to a restricted programming model, in which objects are not allowed to be modified between creation and deletion, and is therefore generally limited. Detlefs et. al. [10] provide a lock-free GC algorithm using reference counting. The approach relies on a strong hardware primitive, namely *double-compare-and-swap* (DCAS) for atomic update of two completely distinct words in memory. Michael [39] presents an efficient lock-free memory management algorithm that does not require special operating system or hardware support. However, his algorithm only guarantees an upper bound on the number of removed nodes not yet freed at any time. This is undesirable because a single garbage node might induce a large amount of occupied resources and might never be reclaimed.

Mark&sweep algorithms do not move objects. They can thus coexist well with C/C++ code, where one never dares to move an object because of possible address computations, and are gaining popularity. Our lock-free mark&sweep algorithm is non-intrusive and features high-performance and reliability. Moreover, unlike most previously published Mark&sweep algorithms [1, 4, 11, 12, 13], we make no assumption on the maximum numbers of mutators and collectors that can operate concurrently. The performance of GC is improved when more processors are involved in it. As far as we could find, no similar algorithm exist.

The correctness properties of any concurrent implementation are seldom easy to verify. This is in general even harder for lock-free algorithms. Our previous work [15, 16] shows that providing correctness proofs for such algorithms require huge amounts of effort, time, and skill. In [17, 18], we have developed two reduction theorems that enable us to reason about a lock-free program to be designed on a higher level than the synchronization primitives *LL/SC* and *CAS*. The reduction theorems are based on refinement mappings as described by Lamport [34], which are used to prove that a lower-level specification correctly implements a higher-level one. Using the reduction theorems, fewer invariants are required and some invariants are easier to discover and formulate without considering the internal structure of the final implementation. In particular, nested loops in the lower-level algorithm may be treated as one loop at a time.

Even so, the structure of our algorithm and its correctness properties, as well as the complexity of reasoning about them, makes neither automatic nor manual verification feasible. We use the higher-

order interactive theorem prover PVS [46] for mechanical support. It is worth noting that there are only a few computer checked correctness proofs of concurrent GC algorithms. In [20], Havelund and Shankar use PVS to verify a safety property of a Mark&sweep GC algorithm, originally suggested by Ben-Ari [4]. In [43], Moreau and Duprat model a distributed reference counting algorithm and prove the safety and liveness property with Coq [9].

Overview of the paper

Section 2 contains the specification of the garbage collector and the interface offered to the users. A higher-level implementation are presented in Section 3. In Section 4, the correctness properties are proven. The proof is based on a list of invariants and lemmas, presented in Appendix A, while the relationships between the invariants are given by a dependency graph in Appendix B. Section 5 gives a low-level lock-free transformation by means of *LL* and *SC* using the reduction theorem developed in [17]. The result is given in Appendix C. Conclusions are drawn in Section 7.

2 Specification

We assume a fixed set `Node` of nodes, each of which is identified with a unique label between 1 and N for some $N \in \mathbb{N}$. To specify garbage collection, we introduce a specification variable `free` of type set of nodes to hold the nodes that are available for allocation of new objects by the application processes. The set `free` is filled by the garbage collectors.

The nodes outside `free` form a finite directed graph of varying structure, called the heap, see Fig. 1. Each node in the graph points to zero or more children (nodes), and the descendent relation may be circular. In the following context, we regard the attributes of nodes as arrays indexed by $1 \dots N$. The number of children of a node x is indicated by its `arity`, which is denoted by `arity[x]` (instead of `Node[x].arity` as it would be if x is an index of array `Node` of some record type with a field `arity`). We let C be the upper bound of the arities of the nodes. The expression `child[x, j]` stands for the pointer to the j th child of node x , where $1 \leq j \leq \text{arity}[x]$.

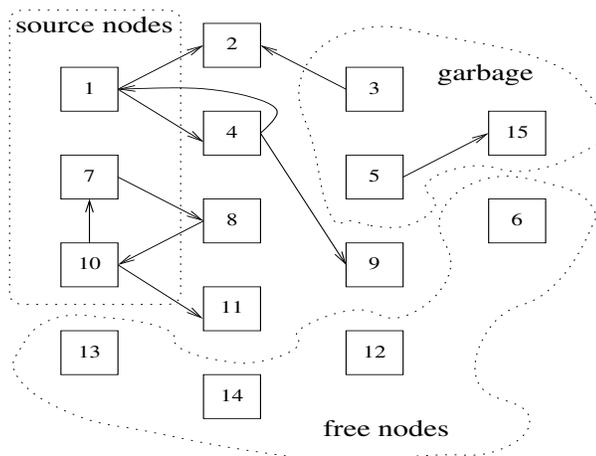


Figure 1: A graph representation of the memory

A node is called a *root* when some process has direct read access to it (such as global/static variables, stack locations and registers of the system). Each application process p maintains a private set `rootsp` that holds its root nodes. The set `Roots` is the union of all `rootsp` for all processes p .

Access to nodes can be transferred between processes. We assume that there is a two-dimensional array `Mbox` indexed with a pair of processes that serves as mailboxes. If process p allows process q to access some node x , it writes x at `Mbox[p, q]` using *Send*. Then, process q can claim access to node x by calling *Receive*.

We call a node a *source node* if the node is either in *Roots* or in some mailbox. A node is called *accessible* iff it is reachable by following a chain of pointers from a source node. Free nodes must not be accessible. Only nodes in the `free` set are allowed to be allocated by the mutators. A node is said to be a *garbage node* if it is neither accessible nor in the `free` set. Behind “user system” garbage collectors compute the set of nodes reachable from a set of source nodes and reclaim all garbage nodes by placing them into the `free` set. More formally, we define

$$\begin{aligned} R(p, x) &\equiv (\exists z \in \text{roots}_p: z \xrightarrow{*} x), \\ R(x) &\equiv (\exists z \in \text{Roots}: z \xrightarrow{*} x) \vee (\exists p, q \in \text{Process}: \text{Mbox}[p, q] \xrightarrow{*} x), \end{aligned}$$

where the reachability relation $\xrightarrow{*}$ is the reflexive transitive closure of relation \rightarrow on nodes defined by

$$z \rightarrow x \equiv (\exists k: 1 \dots \text{arity}[z]: \text{child}[z, k] = x).$$

According to the definitions, a node x is accessible iff $R(x)$ holds. Process p is said to have access to node x if $R(p, x)$ holds. Obviously, $R(p, x)$ implies $R(x)$. The fact that a node x is a garbage node is formalized by:

$$\neg R(x) \wedge x \notin \text{free}.$$

GC does not modify the memory graph (children or arities of nodes) but only repeatedly adds garbage nodes to the `free` set by executing:

```
proc GCCollect()
  < choose  $x \in \text{Node}$  such that  $\neg R(x) \wedge x \notin \text{free}$ ;  $\text{free} := \text{free} \cup x$ ; >
```

Here, and henceforth, we use angular brackets $\langle \rangle$ to indicate that embraced statements are (thought to be) executed atomically.

To specify that GC does happen and is eventually exhaustive, we give the progress property of the collectors specified as follows:

$$\neg R(x) \wedge x \notin \text{free} \Rightarrow \diamond(\neg R(x) \wedge x \in \text{free})$$

that is, every garbage node will be eventually put into the `free` set by a garbage collector.

The machine architecture that we have in mind is based on modern shared-memory multiprocessors that can access a common shared address space. There can be several processes running on a single processor. Let us assume there are P concurrently executing sequential processes along with a set of variables. In the text of a procedure, we use *self* to indicate the process identifier of the process that invokes the procedure. The interface consists of a shared data structure of nodes, and a number of procedures that can be called in the application processes.

We provide procedures that can read and modify the reachable part of the memory graph (from source nodes). An application programmer can assume that the behavior of the routines to access the data is as provided here. In this sense these routines are the specification of our algorithm. In the next chapter we provide implementable routines with the same behavior as specified here. The specification procedures are *Create*, *AddChild*, *GetChild*, *Make*, *Protect*, *UnProtect*, *Send*, *Receive* and *Check*. We use braces $\{ \}$ to indicate a precondition that must hold when invoking a certain procedure.

```
proc Create(): Node
  local  $x$  : Node;
  < when available extract  $x$  from free;
```

```

    arity[x] := 0; rootsself := rootsself ∪ {x};
  return x;

proc AddChild(x, y: Node): Bool
{ R(self, x) ∧ R(self, y) }
  local suc : Bool;
  ⟨ suc := (arity[x] < C);
    if suc then arity[x]++; child[x, arity[x]] := y; fi ⟩
  return suc;

proc GetChild(x: Node, rth: ℕ): Node ∪ {0}
{ R(self, x) }
  local y : Node ∪ {0};
  ⟨ if 1 ≤ rth ≤ arity[x] then y := child[x, rth]; else y := 0; fi ⟩
  return y;

proc Make(c: array [ ] of Node, n: 1..C): Node
{ ∀j: 1 ≤ j ≤ n: R(self, c[j]) }
  local x : Node; j : ℕ;
  ⟨ when available extract x from free;
    for j := 1 to n do child[x, j] := c[j] od;
    arity[x] := n; rootsself := rootsself ∪ {x}; ⟩
  return x;

proc Protect(x: Node)
{ R(self, x) ∧ x ∉ rootsself }
  ⟨ rootsself := rootsself ∪ {x}; ⟩
  return;

proc UnProtect(z: Node)
{ z ∈ rootsself }
  ⟨ rootsself := rootsself \ {z}; ⟩
  return;

proc Send(x: Node, r: Process)
{ R(self, x) ∧ Mbox[self, r] = 0 }
  ⟨ Mbox[self, r] := x; ⟩
  return;

proc Receive(r: Process): Node
{ Mbox[r, self] ≠ 0 }
  local x : Node;
  ⟨ x := Mbox[r, self];
    Mbox[r, self] := 0; rootsself := rootsself ∪ {x}; ⟩
  return x;

proc Check(r, q: Process): Bool
  local suc : Bool;
  ⟨ suc := (Mbox[r, q] = 0); ⟩
  return suc;

```

The application programmers are responsible for ensuring that an offered procedure is called only when its precondition (enclosed by braces if there is any) holds. It is a proof obligation for us that all preconditions of any interface procedure are stable from the perspective of the calling process.

A mutator may continuously allocate a node, add some pointers in the memory, and remove a node from its *roots* set or mailbox. When an allocation request is made, the mutator tries to find a free node (see procedures *Create* and *Make*). The condition “available” in *Create* and *Make* is implementation dependent. When an allocation request cannot be met from the free memory, the mutator either waits, or invokes a new round of GC to free more garbage, or expands the current heap by requesting more memory from the operating system. The threshold value that determines whether or not to invoke a new round of GC can be customized by the user.

The interface is designed in such a way that, when $R(p, x)$ holds, no other process can falsify $R(p, x)$. This means that every process can justify the accessibility of node x by checking $R(p, x)$ (via repeatedly reading arities and children of nodes) without worrying about possible interference from other processes. Indeed, no process is able to decrease $\text{arity}[x]$ or modify $\text{child}[x, j]$ for $1 \leq j \leq \text{arity}[x]$ when node x is accessible (see procedures *AddChild* and *Make*). Instead, the interface only allows to extend the graph by addition of already accessible children. This restriction is stronger than elsewhere, e.g. [4, 32].

The intention of *UnProtect* is that it makes the node and its descendants eligible for garbage collection unless some other process wants to keep them. Via *Send*, *Receive* and *Check*, our algorithm can be used in a distributed system, in which all processors cooperatively traverse the entire data graph by exchanging “messages” to access remote nodes.

3 A higher-level implementation

The idea behind most GC algorithms in use is to first recursively trace all reachable nodes starting from root nodes, then nodes not reached are considered garbage and can be collected. We present a lock-free implementation that comes close to the classical mark&sweep algorithms. Since we allow to transfer access to nodes between processes via mailboxes, we have strengthened the definition of garbage to non-reachability from source nodes instead of from root nodes.

Atomicity is a semantic correctness condition for concurrent systems. Each process in the implementation is a sequential program comprised of labeled groups of statements. Each group is thought to be executed atomically. Actions on private variables can be freely merged to one of the nearest atomic groups without violating the atomicity restriction.

For simplicity, we first extend the specification to a high-level implementation, where all actions on shared variables are separated into distinct atomic accesses except for some special commands enclosed by angular brackets $\langle \dots \rangle$. In order to be able to finally transform the higher-level algorithm into the low-level algorithm using the reduction theorems developed in [17, 18], we require that every labeled atomic group of statements in the higher-level algorithm refer to at most one shared node.

Notational Conventions. Recall that there are P processes with process identifiers ranging from 1 up to P and N nodes labeled from 1 up to N . Unless otherwise specified, we assume that the free variables p , q and r are universally quantified over process identifiers and the free variables w , x , y and z universally quantified over node labels. Since the same sequential program can be executed by all processes, we adopt the convention that every private variable name can be subscripted by the process identifier. In particular, pc_p is the program location of process p . Recall that we regard the attributes of nodes as arrays indexed by $1 \dots N$. E.g. we do not write $\text{Node}[x].f$ but $f[x]$. In order to avoid using too many parentheses, we sometimes use indentation to eliminate brackets and define a binding order for some symbols that we frequently use. The following is a list of these symbols, grouped by binding order; the groups are ordered from the highest binding order to the lowest:

all subscripts and superscripts $()$, $[]$ \neg (not) \wedge (and), \vee (or)

Constant	
P	= number of processes;
N	= number of nodes;
C	= upper bound of number of children;
Type	
colorType	: {white, black, grey};
nodeType	: record =
arity	: \mathbb{N} ; % number of children
child	: array [1...C] of 1...N; % pointers to children
color	: colorType; % holds the color of the node
srcnt	: \mathbb{N} ; % reference counter for a source node
freecnt	: \mathbb{N} ; % dereference counter for a source node
ari	: \mathbb{N} ; % number of children at the beginning of GC
father	: $\mathbb{N} \cup \{-1\}$; % records the parent node GC traverses
round	: \mathbb{N} ; % the latest round of GC involved in, starting from 1
end	
Shared variables	
Node	: array [1...N] of nodeType; % N shared nodes
Mbox	: array [1...P, 1...P] of 0...N; % mailboxes
shRnd	: \mathbb{N} ; % the version of the current round of GC
Private variables	
roots	: a subset of 1...N; % a set of root nodes
rnd	: \mathbb{N} ; % private copy of "shRnd", initially 0!
toBeC	: a subset of 1...N; % a set of nodes to be checked
Initialization:	
shRnd = 1 \wedge $\forall x: 1 \dots N: \text{round}[x] = 1$;	
all other variables are equal to be the minimal values in their respective domains.	

Figure 2: Data Structure

\Rightarrow (implication), \equiv (equivalent) \forall, \exists

3.1 Data Structure

The data structure we use in the higher-level implementation is shown in Fig. 2. We define a shared array `Node` with N elements to represent the memory graph, and a shared two-dimensional array `Mbox` with $P \times P$ elements to represent mailboxes. Besides fields `arity` and `child`, each node has one of three colors: *white*, *black* and *grey*, which is stored in the field `color`. The `free` set is implemented as a virtual set that contains all *white* nodes. All *black* nodes reachable from a source node are interpreted as accessible nodes, and all other *black* nodes are garbage. *Grey* is a transient color that only occurs during GC.

Since any accessible node must not be freed as garbage, the system needs to keep track of source nodes that are created by a process and may still be referred to by other processes. For safety, a process is not allowed to inspect another process's private variable such as `roots`. Instead, we introduce a field `srcnt` for each node to count all references (processes and mailboxes) to the node as a source node. Intentionally, we would like to have something like²:

$$\text{srcnt}[x] = \#\{\{p \mid x \in \text{roots}_p\}\} + \#\{\{(p, q) \mid \text{Mbox}(p, q) = x\}\}.$$

²The precise formula is invariant *I5* in Appendix A.

Therefore, each collector can recognize a source node immediately by checking if its `srcnt` field is positive. We define:

$$R1(x) \equiv (\exists z: \text{srcnt}[z] > 0: z \xrightarrow{*} x),$$

and we have $R(x) \Rightarrow R1(x)$. We do not apply other reference counting to the nodes, since manipulating reference counters is slow and may incur expensive overhead with every duplication and deletion of the pointers.

The main difficulty with tracing the memory graph is that the memory structure can dynamically change during GC. In order to solve this problem, we need some coordination between mutators and collectors to take the view of the memory graph, on which all collectors work. To avoid possible interference between mutators and collectors (we will explain this later), the updates of the field `srcnt` of the node in *UnProtect*, upon deletion from the *roots* set, is postponed until the end of GC. We use the field `freecnt` to count the postponed decrements of `srcnt`. The fields `ari` and `father` contain the number of children a node has at the beginning of GC and the parent node of a node in a tree traversed from a source node by collectors, respectively.

Moreover, since more than one process may participate in GC and they may operate concurrently with mutators, we need to avoid interference from delayed processes. We use a shared variable `shRnd` to hold the round number of the current GC, together with an additional field `round` in the record of a node. The private variable `rnd` is a private copy of the shared variable `shRnd`. A process p participates in the current round of GC if and only if $\text{rnd}_p = \text{shRnd}$. We introduce the global private variable *toBeC* to transfer information about checked nodes between internal calls. There is also a local private variable *toBeD*.

3.2 Algorithm

In this section, we give a higher-level implementation for the collectors and the mutators. Since procedure calls only modify private control data, procedure headers are not always labeled themselves, but their bodies usually have numbered atomic statements. The labels are chosen identical to the labels in the PVS code, and are therefore not completely consecutive.

Brackets `[[]]` and the actions between braces `{ }` and parenthesis `()` can be ignored in the implementation. They only serve in the proof of correctness. We will explain this in section 4.

3.2.1 Collectors

Our garbage collectors are encoded in the procedure *GCollect* as shown in Fig. 3. It is comprised of three phases: (1) paint all *black* nodes *grey* while recording the current memory structure, (2) paint all *grey* nodes reachable from the source nodes back to *black* after traversing the memory graph, and (3) reclaim all garbage by painting all remaining *grey* nodes *white*. The transitions between the colors are shown in Fig. 4.

Processes first let `rnd` get the current value of `shRnd` (this is the only action that updates the private variable `rnd`) to prepare for participating in this round of GC. A new round of GC is started when the fastest process reaches line 101 with $\text{rnd}_{\text{self}} = \text{shRnd}$ holding in the precondition. It is proved by means of invariants that before a new round of GC is started, all earlier rounds of GC have completed:

$$\forall x: 1 \dots N: \text{round}[x] = \text{shRnd} \wedge \text{color}[x] \neq \text{grey}.$$

In order to prevent some process from doing useless or even harmful work, every modification on a node in each phase is protected by a guard, which can force a process (in particular, a process with $\text{rnd}_{\text{self}} \neq \text{shRnd}$) to abandon its delayed operation.

In the first phase, from label 101 to label 108, processes try to update field `round`, paint *black* nodes *grey* and record the present memory structure using fields `ari` and `father`. The processes only need to paint the *black* nodes *grey* since the *white* nodes can not be garbage.

```

proc GCcollect() =
  local  $x: 1 \dots N$ ;  $toBeD$ : a subset of  $1 \dots N$ ;
  % first phase
  100:  $rnd := shRnd$ ;  $toBeC := \{1, \dots, N\}$ ;
  101: while  $shRnd = rnd \wedge toBeC \neq \emptyset$  do
    choose  $x \in toBeC$ ;
  108:   $\langle$  if  $round[x] = rnd$  then
     $round[x] := rnd + 1$ ;  $ari[x] := arity[x]$ ;  $\{ outGC[x] := false; \}$ 
    if  $color[x] = black$  then  $color[x] := grey$ ; fi;
    if  $srcnt[x] > 0$  then  $father[x] := 0$ ; else  $father[x] := -1$ ; fi;  $\rangle$ 
     $toBeC := toBeC \setminus \{x\}$ ;
  od;
  % second phase
  121:  $toBeC := \{1, \dots, N\}$ ;  $toBeD := \{1, \dots, N\}$ ;
  122: while  $shRnd = rnd \wedge toBeD \neq \emptyset$  do
    choose  $x \in toBeD$ ;
  126:   $toBeD := toBeD \setminus \{x\}$ ;
     $\langle$  if  $father[x] = 0$  then  $Mark\_stack(x)$ ; fi;  $\rangle$ 
  od;
  % third phase
  129: while  $shRnd = rnd \wedge toBeC \neq \emptyset$  do
    choose  $x \in toBeC$ ;
  134:   $\langle$  if  $round[x] = rnd + 1 \wedge color[x] = grey$  then
     $color[x] := white$ ;
     $\{ \text{assert } \neg R(x) \wedge x \notin free; free := free \cup x; \}$  fi;  $\rangle$ 
     $toBeC := toBeC \setminus \{x\}$ ;
  od;
  135:  $\langle$  if  $rnd = shRnd$  then  $shRnd := rnd + 1$ ;  $\{ outGC := \lambda(i:1 \dots N):true; \}$  fi;  $\rangle$ 
  137: return
end GCcollect.

```

Figure 3: Procedure *GCcollect*

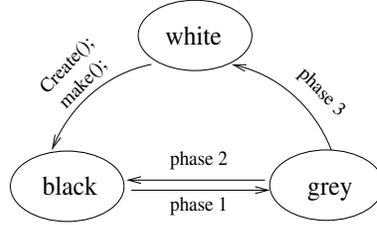


Figure 4: Transitions between these colors.

As the algorithm allows parallel use of mutators, being a source node is not stable during GC. A new source node can be allocated from the `free` set by `Create` or `Make`, or generated by `Protect` or `Send` during GC.

There may be some delay in decrementing field `srcnt` when the number of references decreases (see `UnProtect`). Therefore, we can not say a node x is a source node if its field `srcnt` is positive. Instead, a node is ever a source node since the latest calibration (this is carried out at the end of the second phase) if its field `srcnt` is positive.

We let the field `father` of each node with positive `srcnt` be 0, and that of other nodes be -1 in the first phase. A new source node x can then be distinguished from others by checking if $\text{srcnt}[x] > 0 \wedge \text{father}[x] \neq 0$ holds. For simplicity, we say that a node x with `father` $[x] = 0$ is an *old source node*. When the fastest process participating in the current GC is at the end of its first phase, all non-free nodes are `grey` except that new source nodes are `black`.

A delayed initialization on node x will be skipped because of the guard in line 108 since `round` $[x]$ is never decreased. As usual with version numbers, here we need to assume that sufficient bits are allocated for the version numbers to ensure that they cannot “wrap around” during the interval of a process’s GC cycle.

In the second phase, from label 121 to label 126, the processes build a forest in the set of all reachable nodes starting from the old source nodes. Trees in the forest are mutually disjoint. Each of them is rooted by a chosen old source node, and is created via calling a procedure `Mark_stack` (see Fig. 5) in a `while` loop. During `Mark_stack`, all the `grey` nodes on the tree are painted back to `black` in the order from the leaf to the root.

The procedure `Mark_stack` is mainly a form of graph search, and it was initially designed as a recursive procedure. Since we want to prove the correctness of our algorithm with PVS, we eliminated the recursion in favor of an explicit stack. The private variable `toBeC` serves to ensure that the search of a collector traverses every node at most once. This is important since the memory graph may have cycles and nodes may be reachable from different old source nodes.

In `Mark_stack`, from label 151 to label 163, the tree (in the forest) is established by setting the `father` pointers. Since the memory graph is not a tree and may even have cycles, the processes must reach consensus about the tree. The processes starting from the same old source node cooperate with each other, and are in competition with others to expand the tree to all nodes reached. E.g. in the scenario of Fig. 6, node 7 belongs to tree B since one of the processes establishing tree B first detects that node and sets its `father` to 6. The processes establishing tree A will ignore node 7 since its `father` is now neither -1 nor in tree A . Note that all slower processes starting from the same old source node use the same tree for tracing reachable nodes if the task is not finished. The tree stops growing when every leaf node has no child that regards it as its father.

The order for choosing an element from the local variable `set` is irrelevant for correctness, but relevant for efficiency. The search is a depth first search if the order is first in last out. The search is a breadth first search if the order is first in first out.

The reduction theorems require that every labeled atomic group of statements in the higher-level

```

proc Mark_stack(x: 1...N) =
  local w, y: 1...N; suc: Bool; j, k: ℕ;
  stack: Stack; head: ℕ; set: a subset of 1...N;
  ch: [1...C] of 1...N;
150: toBeC := toBeC \ {x}; set := {x}; head := 0;
151: while shRnd = rnd ∧ set ≠ ∅ do
  choose w ∈ set;
157:   set := set \ {w};
  ⟨ if color[w] = grey ∧ round[w] = rnd + 1 then
    k := ari[w];
    for j := 1 to k do ch[j] := child[w, j] od; ⟩
    head++; stack[head] := w; j := 1;
158:   while shRnd = rnd ∧ j ≤ k do
    y := ch[j];
    if y ∉ toBeC then j++;
    else
163:     ⟨ if (father[y] = -1 ∨ father[y] = w)
      ∧ color[y] = grey ∧ round[y] = rnd + 1 then
        father[y] := w; ⟩
        toBeC := toBeC \ {y}; set := set ∪ {y}; fi;
        j++; fi;
    od; fi;
  od;
168: while shRnd = rnd ∧ head ≠ 0 do
  y := stack[head];
175:   head--;
  ⟨ if color[y] = grey ∧ round[y] = rnd + 1 then
    color[y] := black;
    srcnt[x] := srcnt[x] - freecnt[x]; freecnt[x] := 0; fi; ⟩
  od;
180: return
end Mark_stack.

```

Figure 5: Procedure *Mark_stack*

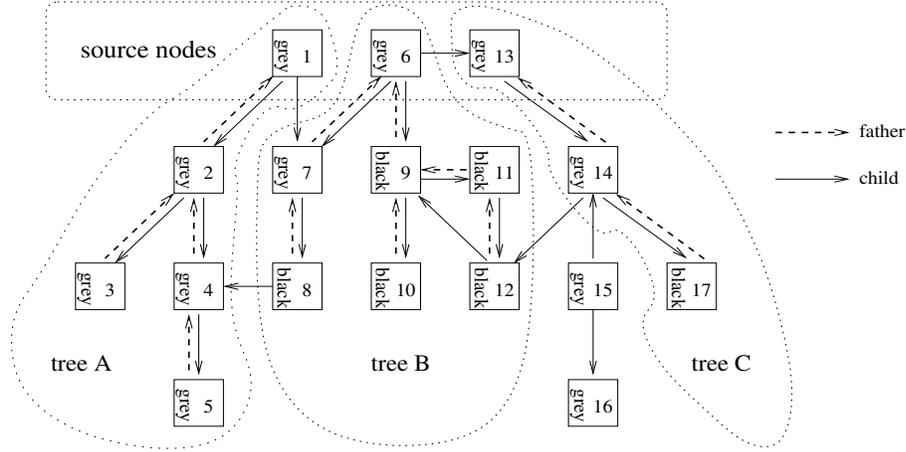


Figure 6: A garbage collection scenario

algorithm refers to at most one shared node. In the procedure *Mark_stack*, local variables *ch* and *k* are therefore introduced to store the old children of a node, temporarily. This also prevents processes from visiting a shared node unnecessarily. It adds a proof obligation that these local variables preserve the information of the node when the process is not delayed.

Starting from the chosen old source node, all nodes on the tree are pushed on the local stack after their old children have been temporarily stored. The order of the elements pushed on the stack is essential for correctness.

After the tree has been established, the process paints all *grey* nodes *black* in the order in which they are popped from the *stack* (from label 168 to label 175) if the action is not too late. When a node in the tree is painted *black*, its descendants (with respect to the *father* relation) in the tree have been painted *black* already (see Fig. 6). So the other processes need not trace or paint the subtree starting from that node. In particular, processes need not trace or paint the tree starting from a new source node. The proof of all this requires interesting and rather complicated graph theoretic invariants. At the end of *Mark_stack*, the process returns to the procedure *GCollect* to search the tree from another old source node.

Note that it is sufficient to explore all accessible *grey* nodes in the second phase without the help of new source nodes. Using the view of the memory structure taken in the first phase may cause to miss collecting some new garbage that is generated by *UnProtecting* a source node after the first phase, but this does not matter since the new garbage will be recycled within two rounds of GC according to the liveness property (we will come to this later).

All old source nodes appear in the different trees of the forest. The tasks of tracing reachable nodes starting from the different old source nodes can be distributed among several processes. When the fastest process that participates in the current GC is at the end of the second phase, all accessible *grey* nodes have been detected and painted *black*.

In the third phase, from label 129 to the end of the procedure, processes try to re-cycle all remaining *grey* nodes by coloring them *white* (i.e. adding them to the *free* set). The main proof obligation for the algorithm is that all nodes being freed are not accessible. When the fastest process that participates in the current GC is at the end of third phase, the shared variable *shRnd* is incremented to notify all other collectors to quit garbage collecting. We define a round of GC to be completed as soon as the fastest process involved in that round of GC finishes the third phase by incrementing *shRnd*.

It is advantageous that the processes may exchange information. The processes involved in the same

round of GC should not use the same strategy for choosing x in the same phase. For the interested reader, more details can be found in the algorithms for the *write-all-problem* [19, 33]. The main idea is to partition the task statically into many roughly equivalent subtasks (more subtasks than the number of available threads), and then let each thread dynamically claim one subtask at a time and remove the subtask after completion.

3.2.2 Mutators

The implementations of the procedures for the mutators are relatively easy. We provide the code in Fig. 7 and Fig. 8 for the interface procedures in the mutators, which match directly with the procedures in the specification. Note that the mutators do not modify fields `ari`, `father` and `round` of nodes.

Procedure *Create* and *Make* serve to extend the memory graph with a new node. In *Create* and *Make*, “time to do GC” indicates that some variable, like time or the amount of free memory, reaches a threshold value. Allocation in the mutator (see *Create* and *Make*) is potentially expensive. It requires a linear search over the whole memory. This problem can be solved by implementing the free set as a lock-free list (see [40, 49]) with adding a new element to the list in a new numbered line just before the last `fi` in line 134, and deletions of elements from the list in lines 200 and 300.

In procedure *UnProtect*, at line 450, the decrementing of the field `srcnt` of the node is postponed when the process removes the node from its *roots* set. Instead, we use the field `freecnt` to count every delayed *UnProtect*. The immediate incrementation of `srcnt` is incorrect because of the following counterexample. Assume there are three nodes: node 1 is a free node, node 2 is a source node, with one child, node 3. Now, process p starts the first phase of GC. Just after process p initializes node 1 (a white node), it goes to sleep. Then process q is scheduled and *Makes* node 1 a new root node, of which the `color` becomes *black* (instead of *grey*), and sets node 3 as a new child of node 1. Then process q *UnProtects* node 2, and node 2 happens to become a non-source node afterwards. Then process p wakes up, resumes to initialize node 2 and node 3. Since in the second phase, processes only explore all *grey* nodes reachable from old source nodes, processes will regard node 3 as an inaccessible node and collect it mistakenly as garbage in the third phase.

One may wonder why the decrementing of the field `srcnt` is postponed from *UnProtect* to line 175 of *Mark_stack*. We tried to update fields `srcnt` in the first phase of GC. However, we found that this is not correct while we proceeded the mechanical proof with PVS. The counterexample is the same as the previous one. After inspecting some invariants, we found that all accessible *grey* nodes can be traced without the help of either the *black* nodes or the upper *grey* nodes resided in the local stack. This means that it is safe to update the field `srcnt` at that moment. Moreover, fields `srcnt` of all remaining *grey* nodes appearing in the third phase are all zero and therefore need not be decremented.

In procedure *Send* and *Receive*, the weaker requirement on the reference counter (i.e. field `srcnt` of a node) is based on the fact that the reference counter does not always need to be accurate.

4 Correctness

The main issue of the algorithm is how to ensure the correct execution of collectors and mutators when they concurrently compete with each other for the same data structure. The standard notion of correctness for asynchronous parallel algorithms is to assume that the atomic instructions of the threads are interleaved in an arbitrary linear order. The algorithm is correct if it behaves properly for all such interleavings. Any property can be considered as the conjunction of safety properties and liveness properties. In this section we describe the proofs of safety properties and a liveness property of the algorithm by means of invariants.

```

proc Create(): 1...N =
  local x: 1...N;
  while true do
200:   choose x ∈ 1...N;
206:   ⟨ if color[x] = white then
        color[x] := black; srcnt[x] := 1;
        ( assert x ∈ free; free := free \ x; )
        [ arity[x] := 0; roots := roots ∪ {x}; ] ⟩
        break;
208:   elseif time to do GC then
        GCollect(); fi;
  od;
210: [ return x ]
end Create.

proc AddChild(x, y: 1...N): Bool =
{ R(self, x) ∧ R(self, y) }
  local suc: Bool;
258: ⟨ [ suc := (arity[x] < C);
        if suc then arity[x]++; child[x, arity[x]] := y; fi ] ⟩
262: [ return suc ]
end AddChild.

proc GetChild(x: 1...N, rth: 1...N): 0...N =
{ R(self, x) }
  local y: 0...N;
280: ⟨ [ if 1 ≤ rth ≤ arity[x] then y := child[x, rth]; else y := 0; fi ] ⟩
284: [ return y ]
end GetChild.

proc Make(c: array [ ] of 1...N, n: 1...C): 1...N =
{ ∀ j: 1...n: R(self, c[j]) }
  local x: 1...N; j: ℕ;
  while true do
300:   choose x ∈ [1...N];
306:   ⟨ if color[x] = white then
        color[x] := black; srcnt[x] := 1;
        ( assert x ∈ free; free := free \ x; )
        [ for j := 1 to n do child[x, j] := c[j]; od
          arity[x] := n; roots := roots ∪ {x}; ] ⟩
        break;
308:   elseif time to do GC then
        GCollect(); fi;
  od;
310: [ return x ]
end Make.

```

Figure 7: Procedure *Create*, *AddChild*, *GetChild* and *Make*

```

proc Protect( $x: 1 \dots N$ ) =
{  $R(\text{self}, x) \wedge x \notin \text{roots}$  }
400:  < srcnt[ $x$ ]++; >
      [[  $\text{roots} := \text{roots} \cup \{x\}$ ; ]]
408:  [[ return ]]
end Protect.

proc UnProtect( $z: 1 \dots N$ ) =
{  $z \in \text{roots}$  }
450:  < freecnt[ $z$ ]++; >
      [[  $\text{roots} := \text{roots} \setminus \{z\}$ ; ]]
460:  [[ return ]]
end UnProtect.

proc Send( $x: 1 \dots N, r: 1 \dots P$ ) =
{  $R(\text{self}, x) \wedge \text{Mbox}[\text{self}, r] = 0$  }
500:  < srcnt[ $x$ ]++; >
508:  [[  $\text{Mbox}[\text{self}, r] := x$ ; ]]
510:  [[ return ]]
end Send.

proc Receive( $r: 1 \dots P$ ):  $1 \dots N$  =
{  $\text{Mbox}[r, \text{self}] \neq 0$  }
  local  $x: 1 \dots N$ ;
550:  [[  $x := \text{Mbox}[r, \text{self}]$ ; ]]
552:  if  $x \notin \text{roots}$  then
      [[  $\text{Mbox}[r, \text{self}] := 0$ ;  $\text{roots} := \text{roots} \cup \{x\}$ ; ]]
  else
558:  < srcnt[ $x$ ]--; >
559:  [[  $\text{Mbox}[r, \text{self}] := 0$ ; ]] ( assert  $x \in \text{roots}$ ; ) fi;
560:  [[ return ]]
end Receive.

proc Check( $r, q: 1 \dots P$ ): Bool
  local  $\text{suc}: \text{Bool}$ ;
600:  [[  $\text{suc} := (\text{Mbox}[r, q] = 0)$ ; ]]
602:  [[ return  $\text{suc}$  ]]
end Check.

```

Figure 8: Procedure *Protect*, *UnProtect*, *Send*, *Receive* and *Check*

4.1 The main loop

In order to verify our memory management system, we model the clients as a very nondeterministic environment in the following loop that may call the interface procedures in any arbitrary order and with arbitrary arguments provided the preconditions are met. This is not part of the memory management system itself, and therefore not to be implemented. It is provided since it is used in the PVS proof to verify the correctness of the system under all possible applications, in the same way as, e.g. in [25] section 4.2.

```

loop
1:  choose call; case call of
      (C) → Create();
      (A, x, y) with R(self, x) ∧ R(self, y) → AddChild(x, y);
      (G, x, rth) with R(self, x) → GetChild(x, rth);
      (M, c, n) with ∀j: 1..n: R(self, c[j]) → Make(c, n);
      (P, x) with R(self, x) ∧ x ∉ rootsself → Protect(x);
      (U, z) with z ∈ rootsself → UnProtect(z);
      (S, x, r) with R(self, x) ∧ Mbox(self, r) = 0 → Send(x, r);
      (R, r) with Mbox(r, self) ≠ 0 → Receive(r);
      (C, r, q) → Check(r, q);
    end
end

```

Normally, after some operation is finished, the process will return to the main loop. In the implementation, there are several places where the same procedure (e.g. *GCCollect*) is called. We introduce an auxiliary private variable *return* to hold the return location. Since they are private, they can be assumed to be touched instantaneously without violation of the atomicity restriction.

4.2 Safety properties

The main aspect of safety is functional correctness and atomicity, say in the sense of [37]. We prove partial correctness of the implementation by showing that each procedure of the implementation executes its specification command always exactly once and that the resulting value of the implementation equals the resulting value in the specification. As shown in Fig. 3 to Fig. 8, we therefore extend the implementations with auxiliary variables and commands used in the specification. For simplicity, we use brackets $\llbracket \ \rrbracket$ to enclose the specification commands that perform the same actions as the implementation, and parenthesis $(\)$ to enclose the specification commands that can be deleted in the implementation.

GC is an internal affair not relevant for the users of the routines. *GCCollect* cannot be invoked explicitly, but will only be invoked implicitly in for instance *Make* and *Create*. This means we only need to prove the match of the specifications and implementations for all user programs, but not for *GCCollect*. Instead, the main safety property we have proved for *GCCollect* is that the system only collects garbage, i.e. that an accessible node is never freed. This is expressed in the invariant:

$$II: \quad \text{color}[x] = \text{white} \Rightarrow \neg R(x).$$

To facilitate the proof of correctness, in *GCCollect* we introduce an auxiliary variable *outGC* to indicate whether a node is not involved in the current round of GC. All operations on *outGC* are enclosed in braces $\{ \ }$, and can be assumed to be executed instantaneously without violation of the atomicity restriction.

We note that the implementation is an extension of the specification except that the set **free** in the specification is implemented as a virtual set of all *white* nodes in the implementation. Apart from the common actions enclosed in $\llbracket \ \rrbracket$, all implementation commands do not modify the specification

variables and all specification commands do not modify the implementation variables. Therefore for simplicity, we do not distinguish the identical variables and commands used in the specification and the implementation, and enclose them in $\llbracket \cdot \rrbracket$. Functional correctness of the mutator now becomes obvious since we have proved the following invariants:

$$\begin{aligned} I2: & \quad \text{color}[x] = \text{white} \equiv x \in \text{free} \\ I3: & \quad 554 \leq pc_p \leq 559 \Rightarrow x_p \in \text{roots}_p \end{aligned}$$

It follows that, by removing the implementation variables from the combined program, we obtain the specification. This removal may eliminate many atomic steps of the implementation. This is known as removal of stutterings in TLA [34] or abstraction from τ steps in process algebras.

Furthermore, we also need to prove that all preconditions of the interface procedures are stable under the actions of the other processes. For the interface procedures *AddChild*, *GetChild*, *Make*, *Protect*, *Send* and *Receive*, the stability of the precondition is expressed respectively by the following invariants:

$$\begin{aligned} I6: & \quad 250 \leq pc_p \leq 258 \Rightarrow R(p, x_p) \wedge R(p, y_p) \\ I7: & \quad pc_p = 280 \Rightarrow R(p, x_p) \\ I8: & \quad 300 \leq pc_p \leq 308 \vee (100 \leq pc_p \leq 180 \wedge \text{return}_p = 300) \Rightarrow \forall k: 1 \dots n_p: R(p, c_p[k]) \\ I9: & \quad pc_p = 400 \vee (500 \leq pc_p \leq 508) \Rightarrow R(p, x_p) \\ I10: & \quad 500 \leq pc_p \leq 508 \Rightarrow \text{Mbox}[p, r_p] = 0 \\ I11: & \quad 550 \leq pc_p \leq 559 \Rightarrow \text{Mbox}[r_p, p] \neq 0 \end{aligned}$$

Process p can ensure its rights to have access to node x by checking the predicate $R(p, x)$, independently, because of the following lemma that asserts $R(p, x)$ can only be invalidated by process p itself:

$$V1: \quad p \neq q \wedge R(p, x) \wedge I18 \wedge I25 \triangleright_q R(p, x),$$

where we write $P \triangleright_q Q$ to express that, if precondition P holds and process q performs an atomic action, this action has postcondition Q .

As we announced earlier, no node is *grey* when the current round of GC is finished. This is formalized in the following invariant:

$$I4: \quad \neg(\exists p: \text{rnd}_p = \text{shRnd}) \Rightarrow \text{color}[x] \neq \text{grey}$$

where the predicate $\text{rnd}_p = \text{shRnd}$ indicates that process p is involved in the current round of GC.

The difference $\text{srcnt}[x] - \text{freecnt}[x]$ of node x counts the number of references to the source node. Since an atomic group of statements in the higher-level implementation must not refer different shared variables (this is an important requirement for the final lock-free transformation), the fields of a node and a mailbox can not be simultaneously modified in the same atomic group. The counter is precisely described by the following invariant:

$$\begin{aligned} I5: \quad \text{srcnt}[x] - \text{freecnt}[x] = & \#(\{p \mid x \in \text{roots}_p\}) + \#(\{(p, q) \mid \\ & (\text{Mbox}(p, q) = x \wedge \neg(pc_q = 559 \wedge p = r_q)) \vee (pc_p = 508 \wedge x_p = x \wedge q = r_p)\}) \end{aligned}$$

All the safety properties (invariants) have been proved with the interactive proof checker PVS. The use of PVS did not only take care of the delicate bookkeeping involved in the proof, it could also deal with many trivial cases automatically. At several occasions where PVS refused to complete the proof, we actually found some mistakes and had to correct previous versions of this algorithm. To prove these invariants, we need many other invariants. All proved invariants and lemmas are listed in Appendix A. Appendix B gives the dependencies between the invariants. For the complete mechanical proof, we refer the interested reader to [24].

4.3 Liveness

A liveness property asserts that program execution eventually reaches some desirable state. In our case, we want to ensure that every garbage node is eventually collected. We shall express this by means of the “leads-to” (denoted as \rightsquigarrow) relation that was developed for UNITY in [7]. For assertions P and Q , we use the formal definition:

$$(P \rightsquigarrow Q) \equiv \Box(P \Rightarrow \Diamond Q).$$

We write $P \triangleright Q$ to express that, if P holds in the precondition of an atomic action, Q holds in the postcondition. Moreover, we define:

$$(P \mathcal{U} Q) \equiv P \triangleright (P \vee Q).$$

All these new symbols are defined to have the same binding order as “ \Rightarrow ” (implication).

The liveness property of the algorithm we need to verify is that, it is always the case that every garbage node is eventually collected. That is,

$$\neg R(x) \rightsquigarrow \text{color}[x] = \text{white}.$$

4.3.1 General Lemmas

In order to prove the liveness property of the algorithm, we establish the needed techniques. First, we introduce fairness into our formalism. This can be done with a single rule: *atomic actions always terminate*. This means that if some process is at the label of some atomic action, this process will eventually execute the action and arrive at the label of the next atomic action. GC is infinitely often triggered during memory allocation when the amount of free memory falls below some threshold or after a certain number of allocations. We therefore need one more fairness assumption, namely GC is eventually called: $\Box(\Diamond(\exists p : pc_p = 100))$.

Except some well-known lemmas extracted from the literatures, all lemmas in this section have been verified mechanically with PVS. The following lemmas are stated in [45].

Lemma 4.1 *For assertions P , Q , R and I ,*

- (a) *Relation \rightsquigarrow is reflexive and transitive.*
- (b) *If $P \Rightarrow Q$ then $P \rightsquigarrow Q$.*
- (c) *If $P \rightsquigarrow R$ and $Q \rightsquigarrow R$ then $(P \vee Q) \rightsquigarrow R$.*
- (d) *If $(P \wedge \Box Q) \rightsquigarrow R$ then $(P \wedge \Box Q) \rightsquigarrow (R \wedge \Box Q)$.*
- (e) *$P \rightsquigarrow (Q \vee (P \wedge \Box \neg Q))$.*

Lemma 4.1 (a), (b) and (c) are used to prove a general *proof lattice* for a program, which is addressed in [45]. Lemma 4.1 (d) shows that in every behavior every state where an invariant holds is followed by states where the invariant always hold. Intuitively, Lemma 4.1 (e) is true because starting from a time where P is true, either Q will be true at some subsequent time, or $\neg Q$ will be always true from then on. Thus, the general pattern of these proofs by contradiction is to assume that the desired predicate never becomes true, and then show that this assumption leads to a contradiction. For more details, refer to [45].

We found the “steps-to” (\triangleright) relation and the “unless” (\mathcal{U}) relation are quite useful to prove the “leads-to” (\rightsquigarrow) relation. Since these two relations only involve a single step, they can be checked directly by PVS with the help of invariants. It is not hard to prove the following general lemmas, which are postulated during the proof of the liveness property.

Lemma 4.2 *For assertions P , Q , R and S ,*

- (a) If P and $(P \rightsquigarrow Q)$ then $\diamond Q$.
- (b) If $P \mathcal{U} Q$ and $\diamond \neg P$ then $P \rightsquigarrow Q$.

Lemma 4.3 Let $Q(w)$ be assertions for all $w \in I$, where I is a finite set. Let P, R, S and T are assertions with $\forall w: I: (P \rightsquigarrow T \vee (S \wedge R \wedge Q(w))) \wedge (S \wedge R \wedge Q(w) \mathcal{U} T \vee \neg S)$ and $\neg S \triangleright \neg S$. Then $P \rightsquigarrow T \vee (S \wedge R \wedge \forall w: I: Q(w))$.

4.3.2 Main Theorems

Actually, we prove something stronger, viz., that, every inaccessible node is painted *white* within two rounds of GC.

Theorem 4.1 For any integer m , $\mathit{shRnd} = m \wedge \neg R(x) \rightsquigarrow \mathit{shRnd} \leq m + 2 \wedge \mathit{white}(x)$.

An invariant has the form of $Q \Rightarrow \square Q$, where Q is an immediate assertion. This means that if the program starts with Q true, then Q is always true throughout its execution. While we proceed the proof of the liveness property, we only need to concern the reachable states starting from initial states where all invariants hold. Therefore, according to Lemma 4.1 (d), we are allowed to add to an assertion any conjunction of invariants freely at any time. To save some space, we denote $\mathit{color}[x] = \mathit{white}$ by $\mathit{white}(x)$, and similarly for the other two colors. Moreover, we define the fastest process that arrives at the first phase, the second phase, the third phase and label 135, and the fastest process that finishes the current GC, respectively by:

$$\begin{aligned}
A_1(m) &\equiv (\exists p : pc_p \in [101, 110] \wedge rnd_p = \mathit{shRnd} = m) \\
&\quad \wedge \neg(\exists p : pc_p \notin [101, 110] \wedge rnd_p = \mathit{shRnd} = m), \\
A_2(m) &\equiv (\exists p : pc_p \notin [101, 110] \wedge rnd_p = \mathit{shRnd} = m) \\
&\quad \wedge \neg(\exists p : pc_p \in [129, 135] \wedge rnd_p = \mathit{shRnd} = m), \\
A_3(m) &\equiv (\exists p : pc_p \in [129, 135] \wedge rnd_p = \mathit{shRnd} = m) \\
&\quad \wedge \neg(\exists p : pc_p = 135 \wedge rnd_p = \mathit{shRnd} = m), \\
A_4(m) &\equiv (\exists p : pc_p = 135 \wedge rnd_p = \mathit{shRnd} = m), \\
A_0(m) &\equiv (\forall p : rnd_p \neq \mathit{shRnd}) \wedge \mathit{shRnd} = m.
\end{aligned}$$

To prove Theorem 4.1, we first prove the following lemmas with PVS, which are related to the “steps-to” (\triangleright) relation and the “unless” (\mathcal{U}) relation.

Lemma 4.4 For any integer m ,

- (a) $A_0(m) \wedge \neg RI(x) \wedge \mathit{black}(x) \wedge \mathit{round}[x] = m$
 $\mathcal{U} A_1(m) \wedge \neg RI(x) \wedge \mathit{black}(x) \wedge \mathit{round}[x] = m$
 $\mathcal{U} A_1(m) \wedge \neg RI(x) \wedge \mathit{grey}(x) \wedge \mathit{round}[x] = m + 1$
 $\mathcal{U} A_2(m) \wedge \neg RI(x) \wedge \mathit{grey}(x) \wedge \mathit{round}[x] = m + 1$
 $\mathcal{U} A_3(m) \wedge \neg RI(x) \wedge \mathit{grey}(x) \wedge \mathit{round}[x] = m + 1$
 $\mathcal{U} A_3(m) \wedge \mathit{white}(x)$
- (b) $\mathit{shRnd} = m = \mathit{round}[x] - 1 \wedge \neg RI(x) \wedge \neg \mathit{white}(x)$
 $\mathcal{U} (\mathit{shRnd} = m \wedge \mathit{white}(x)) \vee (\mathit{shRnd} = m + 1 = \mathit{round}[x] \wedge \neg RI(x) \wedge \neg \mathit{white}(x))$
- (c) $A_0(m) \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \neg R(w) \wedge \mathit{black}(w) \wedge \mathit{srcnt}(w) > 0 \wedge \mathit{round}[w] = m$
 $\mathcal{U} A_1(m) \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \neg R(w) \wedge \mathit{black}(w) \wedge \mathit{srcnt}(w) > 0 \wedge \mathit{round}[w] = m$
 $\mathcal{U} A_1(m) \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \neg R(w) \wedge \mathit{grey}(w) \wedge \mathit{srcnt}(w) > 0 \wedge \mathit{round}[w] = m + 1$
 $\mathcal{U} A_2(m) \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \neg R(w) \wedge \mathit{grey}(w) \wedge \mathit{srcnt}(w) > 0 \wedge \mathit{round}[w] = m + 1$
 $\mathcal{U} A_2(m) \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \mathit{srcnt}(w) = 0$
- (d) $\mathit{shRnd} = m = \mathit{round}[w] - 1 \wedge \neg R(x) \wedge \neg \mathit{white}(x) \wedge \neg R(w) \wedge \mathit{srcnt}(w) > 0$

- $$\begin{aligned} & \mathcal{U} \ (shRnd = m \wedge white(x)) \vee (shRnd = m \wedge \neg R(x) \wedge \neg white(x) \wedge srcnt(w) = 0) \\ & \quad \vee (shRnd = m + 1 = round[w] \wedge \neg R(x) \wedge \neg white(x) \wedge \neg R(w) \wedge srcnt(w) > 0) \\ (e) \quad & shRnd \leq m \wedge \neg R(x) \wedge \neg white(x) \wedge \neg(srcnt(w) > 0 \wedge (w \xrightarrow{*} x)) \\ & \mathcal{U} \ (shRnd \leq m \wedge white(x)) \vee shRnd > m \\ (f) \quad & shRnd > m \triangleright shRnd > m \end{aligned}$$

Using Lemmas 4.1, 4.2, 4.3 and 4.4, we prove the following corollaries.

Corollary 4.1 *For any integer m ,*

$$shRnd = m \rightsquigarrow A_4(m)$$

Proof: Inspired by Lemma 4.1(e), we assume $\Box \neg A_4(m)$. Since the shared variable $shRnd$ can be modified only by some process executing line 135 with precondition $rnd_{self} = shRnd$, we then obtain that $shRnd$ is constant, i.e. m . GC is infinitely often triggered during memory allocation when the amount of free memory falls below some threshold or after a certain number of allocations. We therefore assume that there will eventually exist some process p such that $pc_p = 100$. Because of the fairness of atomic actions, we then get $\Diamond(rnd_p = shRnd = m \wedge pc_p = 101)$. Since $toBeC$ and $toBeD$ are both private variables, all loops in GC are finite (see procedures $GCollect$ and $Mark_stack$). We therefore obtain $\Diamond(rnd_p = shRnd = m \wedge pc_p = 135)$ according to the fairness. This leads to a contradiction. \square

Corollary 4.2 *For any integer m ,*

$$shRnd = m \rightsquigarrow shRnd = m + 1$$

Proof: By Lemma 4.2(a) and Corollary 4.1, we obtain $\Diamond A_4(m)$. Since the shared variable $shRnd$ can be modified only by some process executing line 135 with precondition $rnd_{self} = shRnd$, we then obtain that $shRnd$ will be eventually incremented by 1 according to the fairness. \square

Corollary 4.3 *In Lemma 4.4, all “unless” (\mathcal{U}) relations can be replaced by “leads-to” (\rightsquigarrow) relations.*

Proof: By Lemma 4.2(a) and Corollary 4.2, we obtain $\Diamond(shRnd \neq m)$. Therefore, this corollary is an obvious consequence of using Lemma 4.2(b). \square

Corollary 4.4 *For any integer m ,*

$$shRnd = m = round[x] \wedge \neg R(x) \wedge \neg white(x) \rightsquigarrow shRnd = m \wedge white(x).$$

Proof: By invariant $I16$, we obtain $black(x)$. By invariant $I34$, we have $A_0(m) \vee A_1(m)$. Using transitivity of “leads-to” relation, this is an obvious consequence of Lemma 4.4(a) and Corollary 4.3. \square

Corollary 4.5 *For any integer m ,*

$$shRnd = m \wedge \neg R(x) \wedge \neg R(x) \wedge \neg white(x) \rightsquigarrow shRnd \leq m + 1 \wedge white(x).$$

Proof: By invariant $I13$, we know $round[x] = m \vee round[x] = m + 1$. Therefore, this corollary follows from Corollary 4.4, Lemma 4.4(b) and Corollary 4.3. \square

Corollary 4.6 *For any integer m ,*

$$\begin{aligned} shRnd = m = round[w] \wedge \neg R(x) \wedge \neg white(x) \wedge \neg R(w) \wedge srcnt(w) > 0 & \rightsquigarrow \\ shRnd = m \wedge \neg R(x) \wedge \neg white(x) \wedge srcnt(w) = 0 & \end{aligned}$$

Proof: Obviously, we have $R1(w)$. By invariants *I16* and *I18*, we then obtain $black(w)$. By invariant *I34*, we have $A_0(m) \vee A_1(m)$. Using transitivity of “leads-to” relation, this is an obvious consequence of Lemma 4.4(c) and Corollary 4.3. \square

Corollary 4.7 *For any integer m ,*

$$\begin{aligned} shRnd = m \wedge \neg R(x) \wedge \neg white(x) &\rightsquigarrow (shRnd = m \wedge white(x)) \\ \vee (shRnd \leq m + 1 \wedge \neg R(x) \wedge \neg white(x) \wedge \neg(srcnt(w) > 0 \wedge (w \xrightarrow{*} x))) & \end{aligned}$$

Proof: It holds obviously if $\neg(srcnt(w) > 0 \wedge (w \xrightarrow{*} x))$ is true. Otherwise, by definition of relation R and transitivity of “ $\xrightarrow{*}$ ”, we obtain $\neg R(w)$. By invariant *I13*, we obtain $round[w] = m \vee round[w] = m + 1$. Then, it follows from Corollary 4.6, Lemma 4.4(d) and Corollary 4.3. \square

Corollary 4.8 *For any integer m ,*

$$\begin{aligned} shRnd = m \wedge \neg R(x) \wedge \neg white(x) &\rightsquigarrow \\ (shRnd \leq m + 1 \wedge white(x)) \vee (shRnd \leq m + 1 \wedge \neg R(x) \wedge \neg white(x) \wedge \neg R1(x)). & \end{aligned}$$

Proof: Intending to use Lemma 4.3 with substitutions: $1 \dots N$ for I ; $\neg(srcnt(w) > 0 \wedge (w \xrightarrow{*} x))$ for $Q(w)$; $shRnd = m \wedge \neg R(x) \wedge \neg white(x)$ for P ; $shRnd \leq m + 1$ for S ; $\neg R(x) \wedge \neg white(x)$ for R ; $shRnd \leq m + 1 \wedge white(x)$ for T , it remains to check the premises. The first premise is true because of Corollary 4.7 and Lemma 4.4(e). The second premise is true because of Lemma 4.4(f). \square

Corollary 4.9 *For any integer m ,*

$$shRnd = m \wedge \neg R(x) \wedge \neg white(x) \rightsquigarrow shRnd \leq m + 2 \wedge white(x).$$

Proof: This is an obvious consequence of Lemma 4.1, Corollary 4.5 and Corollary 4.8. \square

Theorem 4.1 follows immediately from Lemma 4.1 and Corollary 4.9.

5 The low-level implementation

Synchronization primitives *LL* and *SC*, proposed by Jensen et al. [29], have found widespread acceptance in modern processor architectures (e.g. MIPS II, PowerPC and Alpha architectures). These instructions are closely related to the *CAS*, and together implement an atomic Read/Write cycle. Instruction *LL* first reads a memory location, say X , and marks it as “reserved” (not “locked”). If no other processor changes the contents of X in between, the subsequent *SC* operation of the same processor succeeds and modifies the value stored; otherwise it fails. There is also a validate instruction *VL*, used to check whether X was not modified since the corresponding *LL* instruction was executed. Implementing *VL* is straightforward in an architecture that already supports *SC*. Note that the implementation does not access or manipulate X by other means than *LL*, *SC* or *VL*. Moir [41] showed that *LL/SC/VL* can be constructed on any system that supports either *LL/SC* or *CAS*.

A shared variable X only accessed by *LL/SC/VL* operations can be regarded as a variable that has an associated shared set of process identifiers $V.X$, which is initially empty. The semantics of *LL*, *SC* and *VL* are given by equivalent atomic statements below.

```

proc LL(ref  $X$ ) : value =
  <  $V.X := V.X \cup \{self\}$ ; return  $X$ ; >

proc SC(ref  $X$ ; in  $Y$ ) : Bool =
  < if  $self \in V.X$  then  $V.X := \emptyset$ ;  $X := Y$ ; return true >

```

```

else return false; fi )

proc VL(ref X) : Bool =
  < return (self ∈ V.X) >

```

5.1 A pattern of general lock-free transformation

At the cost of copying an object’s data before an operation, Herlihy [21] introduced a general methodology to transfer a sequential implementation of any data structure into a lock-free synchronization by means of synchronization primitives *LL* and *SC*. A process that needs access to a shared object pointed by *X* performs a loop of the following steps: (1) read *X* using an *LL* operation to gain access to the object’s data area; (2) make a private copy of the indicated version of the object (this action need not be atomic); (3) perform the desired operation on the private copy to make a new version; (4) finally, call an *SC* operation on *X* to attempt to swing the pointer from the old version to the new version. The *SC* operation will fail when some other process has modified *X* since the *LL* operation, in which case the process has to repeat these steps until consistency is satisfied. The algorithm is non-blocking because at least one out of every *P* attempts must succeed within finite time. Of course, a process might always lose to some faster process, but this is unlikely in practice.

In [17], we formalize Herlihy’s methodology [21] for transferring a sequential implementation of any data structure into a lock-free synchronization using synchronization primitives *LL/SC*, and develop a reduction theorem that enables us to reason about a general lock-free algorithm to be designed on a higher level than the synchronization primitives *LL/SC*.

The lock-free pattern we proposed in [17] is shown in Fig. 9 and Fig. 10, where the following statements are taken as a schematic representation of segments of code:

1. *noncrit*(ref *pub* : aType, *priv* : bType; in *tm* : cType; out *x* : 1 . . . *N*) : representing an atomic non-critical activity on variables *pub* and *priv* according to the value of *tm*, and choosing an index *x* of a shared node to be accessed.
2. *guard*(in *X* : nodeType, *priv* : bType) a non-atomic boolean test on the variable *X* of nodeType. It may depend on private variable *priv*.
3. *com*(ref *X* : nodeType; in *priv* : bType; out *tm* : cType) : a non-atomic action on the variable *X* of nodeType and private variable *tm*. It is allowed to inspect private variable *priv*.
4. *read*(ref *X* : nodeType, in *Y* : nodeType) : a non-atomic read operation that reads the value from the variable *Y* of nodeType to the variable *X* of nodeType, and does nothing else. If *Y* is modified during *read*, the resulting value of *X* is unspecified but type correct, and no error occurs.
5. *LL*, *SC* and *VL*: atomic actions as we defined before.

In the pattern, we are not interested in the internal details of these schematic commands but in their behavior with respect to lock-freedom.

As usual, the action enclosed by angular brackets $\langle \dots \rangle$ is defined as atomic. The private variable *x* is intended only to determine the node under consideration, the private variable *tm* is intended to hold the result of the critical computation *com*, if executed.

In the concrete system \mathcal{S}_c of Fig. 10, we declare *P* extra shared nodes for private use (one for each process). Array *indir* acts as pointers to shared nodes. *node*[*mp*_{*p*}] can always be taken as a “private” node of process *p* though it is declared publicly: other processes can read it but cannot modify it. If some other process successfully updates a shared node while an active process *p* is copying the shared node to its “private” node, process *p* will restart the inner loop, since its private view of the node is not consistent anymore. After the assignment *mp* := *m* at line *c6*, the “private” node becomes

```

Constant
  P = number of processes;
  N = number of nodes;
Shared variable
  pub: aType;
  Node: array [1...N] of nodeType;
Private variable
  priv: bType; pc: {a1, a2}; x: 1...N; tm: cType;

Program
  loop
a1:  noncrit(pub, priv, tm, x);
a2:  ( if guard(Node[x], priv) then com(Node[x], priv, tm); fi; )
  end.

Initial conditions
   $\Theta_a: \forall p: 1 \dots P: pc_p = a1$ 
Liveness
   $\mathcal{L}_a: \Box(pc_p = a2 \longrightarrow \Diamond pc_p = a1)$ 

```

Figure 9: Interface \mathcal{S}_a

```

Constant
  P = number of processes;
  N = number of nodes;
Shared variable
  pub: aType;
  node: array [1...N+P] of nodeType;
  indir: array [1...N] of 1...N+P;
Private variable
  priv: bType; pc: [c1...c7];
  x: 1...N; mp, m: 1...N+P; tm, tm1: cType;

Program
  loop
c1:  noncrit(pub, priv, tm, x);
  loop
c2:  m := LL(indir[x]);
c3:  read(node[mp], node[m]);
c4:  if guard(node[mp], priv) then
c5:    com(node[mp], priv, tm1);
c6:    if SC(indir[x], mp) then
      mp := m; tm := tm1; break; fi;
c7:  elseif VL(indir[x]) then break; fi; fi;
  end;
  end.

Initial conditions
   $\Theta_c: (\forall p: 1 \dots P: pc_p = c1 \wedge mp_p = N+p) \wedge (\forall i: 1 \dots N: indir[i] = i)$ 
Liveness
   $\mathcal{L}_c: \Box(pc_p = c2 \longrightarrow \Diamond pc_p = c1)$ 

```

Figure 10: Lock-free implementation \mathcal{S}_c of \mathcal{S}_a

shared and the node shared previously (which contains the old version) becomes “private”. Keep in mind that the composition of `node` and `indir` in \mathcal{S}_c corresponds to `Node` in the abstract system \mathcal{S}_a of Fig. 9.

The following theorem is the reduction theorem stated in [17].

Theorem 5.1 *The abstract system \mathcal{S}_a defined in Fig. 9 is refined by the concrete system \mathcal{S}_c defined in Fig. 10, that is, $\exists\phi: \mathcal{S}_c \sqsubseteq \mathcal{S}_a$.*

The reduction theorem is based on refinement mapping as described by Lamport [34], which asserts that a lower-level specification correctly implements a higher-level one. It has been verified with PVS. A reduction theorem is a general rule for deriving an “equivalent” higher-level specification from a lower-level one in some suitable sense [8].

5.2 The lock-free implementation

Refinement mappings enable us to reduce an implementation by reducing its components in relative isolation, and then gluing the *reductions* together with the same structure as the implementation. Atomicity guarantees that a parallel execution of a program gives the same results as a sequential and non-deterministic execution. This allows us to use the refinement calculus for stepwise refinement of transition systems [2]. Therefore, Theorem 5.1 can be universally employed for a lock-free construction to synchronize access to shared nodes of `nodeType`, and be sure that we end up with the reduction of the implementation. This allows us to design and verify a lock-free program on a higher level than the synchronization primitives. The big advantage is that substantial pieces of the concrete program can be dealt with as atomic statements on the higher level and thus the correctness can be more easily verified.

In the higher-level implementation (from Fig. 3 to Fig. 8), instruction 135 is simply a *CAS* instruction offered by machine architectures or a Read/Write cycle that can easily be implemented by a *LL/SC*. Each of all other special commands enclosed by angular brackets $\langle \dots \rangle$ only refer one shared node and some private variables, and therefore can be transformed into low-level lock-free implementations using Theorem 5.1. E.g. line 108 of Fig. 3 is implemented in lines 104... 107 of Appendix C.2, where `round[mp]` is an alias of `Node[mp].round`, and similarly for `color[mp]`, `srcnt[mp]` and `father[mp]`. At lines 126 and 157 (and possibly other cases), since these commands do not modify the node, swapping of pointers is unnecessary. We therefore use a simplified version where *SC* can be replaced by *VL*. After the transformation, all statements in the algorithm are atomic³. The transformation is straightforward, and we present our final lock-free algorithm in Appendix C.

Apart from that, the higher-level algorithm can also be transformed into a lock-free implementation by means of *CAS* using the other reduction theorem developed in [18]. This final transformation is a bit more complicated. Because of the similarity, we don’t provide the theorem and the final transformation here.

6 Practical experiment

We carried out a number of experiments with our algorithm in order to obtain some insight in its practical performance. The first major conclusion is that the performance of the algorithm is strongly influenced by its parameter setting such as the total number of nodes, the condition for joining the garbage collection process, the percentage of occupied nodes and the division of work between garbage collecting and manipulating the data structure. A second conclusion is that if we set these parameters well, we see no degrading in performance when increasing the number of processes. A third conclusion

³Note that accesses to “private” nodes do not violate the atomicity restriction and can be freely added to an atomic statement.

is that due to the fact that garbage collection is a relatively elaborate affair, the performance in terms of the numbers of nodes that are created and collected per unit of time is relatively low.

#Processes	1 processor		2 processors		4 processors	
	0	100	0	100	0	100
1	833k	125k	800k	87k	714k	77k
2	800k	118k	784k	200k	750k	136k
3	789k	115k	741k	207k	794k	207k
4	800k	114k	727k	205k	800k	216k
5	794k	114k	758k	204k	800k	233k
10	833k	112k	870k	222k	851k	263k
15	789k	115k	833k	214k	845k	261k
20	769k	118k	769k	211k	800k	258k
25	781k	114k	735k	208k	840k	256k
31	756k	115k	729k	214k	844k	253k

Table 1: Some experimental results

In our experimental setup, we let a number of processes repeatedly create a node, read it a number of times and release it again. One process is the garbage collector process, and the settings of parameters are such that no other process will join in to assist this process. However, if a process fails too often to obtain a free node, the process yields its processor, putting itself in the processor waiting queue. This provides an effective load balancing policy. Note that in order to maintain the lock-free nature of the algorithm this process must eventually join garbage collection if obtaining free nodes fails continuously.

More concretely, in the experiments reported in table 6, we use a small array of 2000 nodes and let each process create a large ($> 10^6$) number of nodes. Each processor that must create a node tries 15 times to find a free node before yielding its processor.

The experiments have been carried out on a one, two and four processor machine. All machines were Intel Linux machines of the following types:

- The single processor uses a Pentium III (Coppermine) processors of 1Ghz with 256kb cache each.
- The two processor machine contained two Xeon CPUs of 2.8Ghz with 512kb of cache each.
- The four processor machine has four Intel Xeon CPUs of 2.4 Ghz with 512 KB cache each.

The load linked, store conditional and verify link statements have been implemented using the 64 bit compare and swap (cmpxchg8B) instruction available on Intel Pentium processors (see [42] for the implementation). This limits parallelism to 32 processes, but does not have problems with wrap around as the implementation in [41].

The table provides the number of nodes that could be created and read per second. The letter ‘k’ indicates that the figures refer to thousands of nodes. In the columns marked with 0 these nodes are read 0 times, and in the columns marked with 100 these nodes are read a 100 times. The column headed with #Processes indicates the number of processes that were used to create new nodes. As stated above there is one additional process doing garbage collecting.

Note that the table shows an almost perfect linear scaling. The variations in the table can fully be contributed to statistical noise. Only when there are few processes on a multiprocessor machine performance is bad. This is due to the fact that the garbage collector has plenty of time compared to the heavily loaded processes, – which must read often – and therefore wastes its time. Note also that on multiprocessor machines the performance on generating nodes is low compared to the relatively slow single processor machine. We believe that this is due to interprocess communication induced by compare and swap.

7 Conclusions

We present a lock-free parallel algorithm for mark&sweep GC in a realistic model by means of synchronization primitives *load-linked (LL)/store-conditional (SC)* or *CAS* offered by machine architectures. Our algorithm allows to collect a circular data structure and makes no assumption on the maximum number of mutators and collectors that can operate concurrently during GC. The efficiency of GC can be enhanced when more processors are involved in it. Providing *Send* and *Receive*, our algorithm can be adapted to a distributed system, in which all processors cooperatively traverse the entire data graph by exchanging “messages” to access remote nodes.

Formal verification is desirable because there could be subtle bugs as the complexity of algorithms increases. To ensure our correctness proof presented in the paper is not flawed, we use the higher-order interactive theorem prover PVS for mechanical support. PVS has a convenient specification language and contains a proof checker which allows users to construct proofs interactively, to automatically execute trivial proofs, and to check these proofs mechanically. At several occasions where PVS refused to let a proof be finished, we actually found a mistake and had to correct previous versions of the algorithm. For the complete mechanical proof, we refer the reader to [24].

The entrenched problem inherited from classical mark&sweep algorithms is that our algorithm may also result in severe memory fragmentation, with lots of small blocks. It is possible that there will be no block of memory on the free list large enough to hold a large object, such as an array. Thus, it is important to move free blocks that happen to be adjacent in memory. We plan in the future to incorporate some appropriate copying technique in our algorithm.

A Invariants

We present here all invariants and lemmas whose validity has been verified with the theorem prover PVS. In the invariants and the lemmas below, we use the relations $R(x)$, $R(p, x)$, $R1(p, x)$, \triangleright_q defined in sections 2, 2, 3.1 and 4.2, respectively. The relation $\xrightarrow{*}$ is defined in section 2. The relation $\xrightarrow{M*}$ is the reflexive transitive closure of relation \xrightarrow{M} on nodes defined by:

$$z \xrightarrow{M} x \equiv (\text{color}[z] = \text{black} \wedge \text{aux}[z] \wedge \exists k: 1 \dots \text{arity}[z]: \text{child}[z, k] = x) \\ \vee (\text{color}[z] = \text{grey} \wedge \exists k: 1 \dots \text{ari}[z]: \text{child}[z, k] = x)$$

We define the j -th ancestor of a node x by the recursive function:

$$\text{anc}(x, j) \equiv ((j = 0 \vee \text{father}[x] \leq 0) ? x : \text{anc}(\text{father}[x], j - 1))$$

Main invariants:

- I1: $\text{color}[x] = \text{white} \Rightarrow \neg R(x)$
- I2: $\text{color}[x] = \text{white} \equiv x \in \text{free}$
- I3: $554 \leq pc_p \leq 559 \Rightarrow x_p \in \text{roots}_p$
- I4: $\neg(\exists p: \text{rnd}_p = \text{shRnd}) \Rightarrow \text{color}[x] \neq \text{grey}$
- I5: $\text{srcnt}[x] - \text{freecnt}[x] = \#\{\{p \mid x \in \text{roots}_p\}\} + \#\{\{p, q \mid \\ (\text{Mbox}(p, q) = x \wedge \neg(pc_q = 559 \wedge p = r_q)) \vee (pc_p = 508 \wedge x_p = x \wedge q = r_p)\}\}$

Invariants about the stability of the preconditions in the offered procedures:

- I6: $250 \leq pc_p \leq 258 \Rightarrow R(p, x_p) \wedge R(p, y_p)$
- I7: $pc_p = 280 \Rightarrow R(p, x_p)$
- I8: $300 \leq pc_p \leq 308 \vee (100 \leq pc_p \leq 180 \wedge \text{return}_p = 300) \Rightarrow \forall k: 1 \dots n_p: R(p, c_p[k])$
- I9: $pc_p = 400 \vee (500 \leq pc_p \leq 508) \Rightarrow R(p, x_p)$
- I10: $500 \leq pc_p \leq 508 \Rightarrow \text{Mbox}[p, r_p] = 0$
- I11: $550 \leq pc_p \leq 559 \Rightarrow \text{Mbox}[r_p, p] \neq 0$

Invariants that hold globally:

- I12: $rnd_p \leq shRnd$
- I13: $shRnd \leq round[x] \leq shRnd + 1$
- I14: $\neg aux[x] \Rightarrow round[x] = shRnd + 1$
- I15: $\neg(\exists p: rnd_p = shRnd) \Rightarrow round[x] \leq shRnd$
- I16: $round[x] \leq shRnd \Rightarrow color[x] \neq grey$
- I17: $color[x] = grey \Rightarrow \neg aux[x]$
- I18: $color[x] = white \Rightarrow \neg RI(x)$
- I19: $color[x] = white \Rightarrow father[x] \leq -1$
- I20: $color[x] = grey \vee father[x] \geq 0 \Rightarrow ari[x] \leq arity[x]$
- I21: $color[x] = grey \wedge father[x] > 0 \Rightarrow \exists k: 1 \dots ari[father[x]]: child[father[x], k] = x$
- I22: $x \neq father[x]$
- I23: $father[x] > 0 \Rightarrow \neg(\exists j: \mathbb{N}: x = anc(x, j))$
- I24: $father[x] = 0 \wedge color[x] = grey \Rightarrow srcnt[x] > 0$
- I25: $(\exists p: x \in roots_p) \Rightarrow srcnt[x] - freecnt[x] > 0$
- I26: $\neg R(x) \Rightarrow srcnt[x] - freecnt[x] = 0$
- I27: $RI(x) \wedge (color[x] = grey \vee (color[x] = black \wedge aux[x]))$
 $\Rightarrow \exists w: srcnt[w] > 0 \wedge (father[w] = 0 \vee aux[w]) \wedge w \xrightarrow{M^*}$
- I28: $return_p = 200 \vee return_p = 300 \vee return_p = 450$

Invariants about the first phase of GC:

- I29: $101 \leq pc_p \leq 110 \wedge rnd_p = shRnd \wedge \neg(x \in toBeC_p \wedge aux[x])$
 $\Rightarrow round[x] = rnd_p + 1$
- I30: $101 \leq pc_p \leq 110 \wedge rnd_p = shRnd \wedge \neg(\exists r: \neg(101 \leq pc_r \leq 110) \wedge rnd_r = shRnd)$
 $\Rightarrow father[x] \neq 0 \vee aux[x] \vee color[x] \neq black$
- I31: $101 \leq pc_p \leq 110 \wedge rnd_p = shRnd \wedge \neg(\exists r: \neg(101 \leq pc_r \leq 110) \wedge rnd_r = shRnd)$
 $\wedge (x \notin toBeC_p \vee round[x] = rnd_p + 1)$
 $\Rightarrow \neg aux[x]$
- I32: $101 \leq pc_p \leq 110 \wedge rnd_p = shRnd \wedge \neg(\exists r: \neg(101 \leq pc_r \leq 110) \wedge rnd_r = shRnd)$
 $\wedge \neg(x \in toBeC_p \wedge aux[x])$
 $\Rightarrow father[x] = 0 \vee father[x] = -1$
- I33: $101 \leq pc_p \leq 110 \wedge rnd_p = shRnd \wedge \neg(\exists r: \neg(101 \leq pc_r \leq 110) \wedge rnd_r = shRnd)$
 $\wedge (x \notin toBeC_p \vee round[x] = rnd_p + 1) \wedge srcnt[x] = 0$
 $\Rightarrow father[x] = -1$

Invariants about the second phase of GC:

- I34: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \Rightarrow round[x] = rnd_p + 1$
- I35: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \Rightarrow \neg aux[x]$
- I36: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \wedge father[w] \leq -1 \Rightarrow \neg(\exists x: father[x] = w)$
- I37: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \wedge color[x] = grey \wedge father[x] > 0$
 $\Rightarrow color[father[x]] = grey \wedge father[father[x]] \geq 0$
- I38: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \wedge father[x] > 0 \wedge color[x] = grey$
 $\Rightarrow \exists j: 1 \dots N: father[anc(x, j)] = 0 \wedge color[anc(x, j)] = grey$
- I39: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \wedge \neg RI(x) \wedge color[x] = grey$
 $\Rightarrow father[x] = -1$
- I40: $\neg(101 \leq pc_p \leq 110) \wedge rnd_p = shRnd \wedge color[w] = black \wedge father[w] \geq 0$
 $\Rightarrow \forall k: 1 \dots ari[w]: (father[child[w], k]) = w \vee father[child[w], k] < 0$
 $\Rightarrow color[child[w], k] = black$
- I41: $pc_p = 121 \Rightarrow rnd_p \neq shRnd \vee toBeC_p = \emptyset$
- I42: $\neg(101 \leq pc_p \leq 121) \wedge rnd_p = shRnd \wedge srcnt[x] > 0 \wedge father[x] = 0$
 $\wedge \neg(x \in toBeD_p \vee (150 \leq pc_p \leq 180 \wedge x = x_p))$
 $\Rightarrow color[x] = black$

- I43: $(122 \leq pc_p \leq 127 \vee 150 \leq pc_p \leq 180) \wedge rnd_p = shRnd \wedge x \notin toBeC_p$
 $\Rightarrow father[x] \geq 0$
- I44: $((122 \leq pc_p \leq 127 \vee 150 \leq pc_p \leq 180) \wedge rnd_p = shRnd \wedge x \in toBeD_p$
 $\wedge father[x] = 0) \vee (pc_p = 150 \wedge rnd_p = shRnd \wedge x = x_p)$
 $\Rightarrow x \in toBeC_p$
- I45: $(122 \leq pc_p \leq 127 \vee 150 \leq pc_p \leq 180) \wedge rnd_p = shRnd \wedge x \notin toBeC_p$
 $\Rightarrow \neg(\exists w: father[x] = w \wedge (w \in set_p \vee w \in toBeC_p))$
- I46: $122 \leq pc_p \leq 180 \wedge rnd_p = shRnd$
 $\wedge \neg(x \in toBeC_p \vee (pc_p \geq 151 \wedge (x \in set_p \vee \exists i: 1 \dots head_p: x = stack_p[i])))$
 $\Rightarrow color[x] \neq grey$

Invariants about procedure *Mark_stack*:

- I47: $pc_p = 150 \Rightarrow x_p \notin toBeD_p$
- I48: $150 \leq pc_p \leq 180 \wedge rnd_p = shRnd$
 $\Rightarrow (color[x_p] = grey \wedge father[x_p] = 0 \wedge srcnt[x_p] > 0)$
 $\vee (color[x_p] = black \wedge father[x_p] = 0)$
- I49: $151 \leq pc_p \leq 180 \wedge x \in toBeC_p$
 $\Rightarrow \neg(x \in set_p \vee \exists i: 1 \dots head_p: x = stack_p[i])$
- I50: $151 \leq pc_p \leq 180 \wedge (\exists i: 1 \dots head_p: x = stack_p[i])$
 $\Rightarrow x \notin set_p \wedge x \notin toBeC_p$
- I51: $151 \leq pc_p \leq 180 \wedge rnd_p = shRnd$
 $\Rightarrow x_p \in set_p \vee color[x_p] = black \vee (\exists i: 1 \dots head_p: x_p = stack_p[i])$
- I52: $151 \leq pc_p \leq 180 \wedge rnd_p = shRnd \wedge (x \in set_p \vee \exists i: 1 \dots head_p: x = stack_p[i])$
 $\Rightarrow (color[x] = grey \vee color[x] = black) \wedge father[x] \geq 0$
- I53: $151 \leq pc_p \leq 180 \wedge rnd_p = shRnd \wedge color[x] = grey$
 $\wedge (x \in set_p \vee \exists i: 1 \dots head_p: x = stack_p[i])$
 $\Rightarrow RI(x) \wedge father[x] \geq 0$
- I54: $151 \leq pc_p \leq 180 \wedge rnd_p = shRnd \wedge (\exists i: 1 \dots head_p: w = stack_p[i])$
 $\Rightarrow \forall k: 1 \dots ari[w]: (father[child[w, k]] \geq 0 \vee color[child[w, k]] = black)$
 $\wedge (father[child[w, k]] = w \Rightarrow child[w, k] \in set_p \vee color[child[w, k]] = black$
 $\vee (\exists j: 1 \dots head_p: child[w, k] = stack_p[j]) \wedge (\forall m, n: 1 \dots head_p:$
 $w = stack_p[m] \wedge x = stack_p[n] \Rightarrow m < n))$
 $\vee (158 \leq pc_p \leq 164 \wedge w_p = w \wedge k \geq j_p)$
- I55: $pc_p = 158 \wedge rnd_p = shRnd \Rightarrow j_p = 1 \vee 1 < j_p \leq ari[w_p] + 1$
- I56: $158 \leq pc_p \leq 164 \wedge rnd_p = shRnd$
 $\Rightarrow k_p = ari[w_p] \wedge \forall j: 1 \dots k_p: ch_p[j] = child[w_p, j]$
- I57: $158 \leq pc_p \leq 164 \wedge rnd_p = shRnd \wedge \neg(x \in toBeC_p \vee \exists j: 1 \dots j_p - 1: x = child[w_p, j])$
 $\Rightarrow father[x] \neq w_p$
- I58: $158 \leq pc_p \leq 164 \wedge rnd_p = shRnd$
 $\Rightarrow \forall k: 1 \dots j_p - 1: (color[child[w_p, k]] = grey \Rightarrow father[child[w_p, k]] \geq 0$
 $\wedge (father[child[w_p, k]] = w_p \Rightarrow child[w_p, k] \in set_p))$
- I59: $158 \leq pc_p \leq 165 \Rightarrow \exists i: 1 \dots head_p: w_p = stack_p[i]$
- I60: $159 \leq pc_p \leq 164 \wedge rnd_p = shRnd \Rightarrow 1 \leq j_p \leq ari[w_p] \wedge y_p = child[w_p, j_p]$
- I61: $168 \leq pc_p \leq 180 \Rightarrow rnd_p \neq shRnd \vee set_p = \emptyset$
- I62: $170 \leq pc_p \leq 176 \Rightarrow head_p \neq 0$
- I63: $pc_p = 180 \Rightarrow rnd_p \neq shRnd \vee head_p = 0$

Invariants about the third phase of GC:

- I64: $129 \leq pc_p \leq 137 \wedge rnd_p = shRnd \wedge color[x] = grey \Rightarrow \neg RI(x)$
- I65: $pc_p = 134 \wedge round[x] = rnd_p + 1 \wedge color[x] = grey$
 $\Rightarrow \neg R(x) \wedge x \notin free$
- I66: $\neg(101 \leq pc_p \leq 134 \vee 150 \leq pc_p \leq 180) \wedge rnd_p = shRnd$
 $\Rightarrow color[x] \neq grey$

I67: $pc_p = 135 \Rightarrow rnd_p \neq shRnd \vee toBeC_p = \emptyset$
I68: $\neg(101 \leq pc_p \leq 135 \vee 150 \leq pc_p \leq 180) \Rightarrow rnd_p \neq shRnd$

Invariants outside GC:

I69: $pc_p = 450 \vee (100 \leq pc_p \leq 180 \wedge return_p = 450) \Rightarrow R(p, z_p)$
I70: $500 \leq pc_p \leq 508 \Rightarrow Mbox[p, r_p] = 0$
I71: $552 \leq pc_p \leq 559 \Rightarrow x_p = Mbox[r_p, p] \wedge x_p \neq 0$
I72: $pc_p = 558 \Rightarrow srcnt[x_p] > 1$

Main lemmas:

V1: $p \neq q \wedge R(p, x) \wedge I18 \wedge I25 \triangleright_q R(p, x)$
V2: $color[x] \neq white \wedge \neg R1(x) \wedge I6 \wedge I8 \wedge I9 \wedge I25 \wedge I63 \wedge I66 \triangleright \neg R1(x)$

B Dependencies between invariants

Let us write “ ϕ **from** ψ_1, \dots, ψ_n ” to denote that ϕ can be proved to be an invariant using that ψ_1, \dots, ψ_n hold in the precondition of every step. We write “ $\phi \Leftarrow \psi_1, \dots, \psi_n$ ” to denote that ϕ can be directly derived from ψ_1, \dots, ψ_n . We have verified the following “**from**” and “ \Leftarrow ” relations mechanically:

I1 \Leftarrow I3, I5, I18, I71
I2 **from** : *true*
I3 **from** : I28
I4 \Leftarrow I12, I15, I16
I5 **from** : I18, I25, I28, I70, I71
I6 **from** : I18, I25, I28
I7 **from** : I18, I25, I28
I8 **from** : I18, I25, I28
I9 **from** : I18, I25, I28
I10 **from** : I28
I11 **from** : I28
I12 **from** : *true*
I13 **from** : I12, I34
I14 **from** : I12, I13
I15 **from** : I12, I13, I34
I16 **from** : I12, I13, I66
I17 **from** : I66
I18 **from** : I6, I8, I9, I12, I16, I25, I64, I69
I19 **from** : I12, I16, I18, I39, I64
I20 **from** : I19
I21 **from** : I12, I13, I14, I15, I17, I20, I32, I34, I37, I60, I66
I22 \Leftarrow I23
I23 **from** : I13, I16, I36, I50, I59
I24 **from** : *true*
I25 \Leftarrow I5
I26 \Leftarrow I3, I5, I9
I27 **from** : I6, I8, I9, I12, I13, I14, I16, I17, I18, I20, I21, I24, I25, I35, I37, I38, I54, I61, I64, I66, I69
I28 **from** : *true*
I29 **from** : I12, I13, I14, I28
I30 **from** : I12, I13, I14, I15, I16, I19, I28, I68
I31 **from** : I12, I13, I15, I28, I35, I68
I32 **from** : I12, I13, I14, I15, I28, I29, I31, I34, I68


```

    srcnt:  $\mathbb{N}$ ; % reference counter for a source node
    freecnt:  $\mathbb{N}$ ; % dereference counter for a source node
    ari:  $\mathbb{N}$ ; % number of children at the beginning of GC
    father:  $\mathbb{N} \cup \{-1\}$ ; % records the parent node GC traverses
    round:  $\mathbb{N}$ ; % the latest round of GC involved in, starting from 1
end
Shared variables
node: array [1...N + P] of nodeType; % N shared nodes, P implicitly private nodes
indir: array [1...N] of 1...N + P; % N pointers to N shared nodes
Mbox: array [1...P, 1...P] of 0...N; % mailboxes
shRnd:  $\mathbb{N}$ ; % the version of the current round of GC
Private variables
roots: a subset of 1...N; % a set of root nodes
rnd:  $\mathbb{N}$ ; % private copy of "shRnd", initially 0!
toBeC: a subset of 1...N; % a set of nodes to be checked
mp: 1...N + P; % the pointer to the private copy of a node

Initialization:
shRnd = 1  $\wedge$   $\forall x: 1 \dots N: (\text{indir}[x] = x \wedge \text{round}[\text{indir}[x]] = 1)$ ;
 $\forall p: 1 \dots P: mp_p = N + p$ ;
all other variables are equal to be the minimal values in their respective domains.

```

C.2 Algorithm

```

proc GCollect() =
    local m : 1...N + P; x: 1...N; toBeD: a subset of 1...N;
% first phase
100: rnd := shRnd; toBeC := {1, ..., N};
101: while shRnd = rnd  $\wedge$  toBeC  $\neq$   $\emptyset$  do
    choose  $x \in$  toBeC;
    while true do
102:     m := LL(indir[x]);
103:     node[mp] := node[m];
104:     if round[mp] = rnd then
105:         round[mp] := rnd + 1; ari[mp] := arity[mp];
            if color[mp] = black then color[mp] := grey; fi;
            if srcnt[mp] > 0 then father[mp] := 0; else father[mp] := -1; fi;
106:         if SC(indir[x], mp) then toBeC := toBeC - {x}; mp := m; break; fi;
107:         elseif VL(indir[x]) then toBeC := toBeC - {x}; break; fi;
    od;
od;
% second phase
110: toBeC := {1, ..., N}; toBeD := {1, ..., N};
111: while shRnd = rnd  $\wedge$  toBeD  $\neq$   $\emptyset$  do
    choose  $x \in$  toBeD;
    while true do
112:     m := LL(indir[x]);
113:     node[mp] := node[m];
114:     if father[mp] = 0 then
116:         if VL(indir[x]) then
            toBeD := toBeD - {x};
            Mark_stack(x); break; fi;
117:         elseif VL(indir[x]) then toBeD := toBeD - {x}; break; fi;
    od;
od;

```

```

% last phase
120: while shRnd = rnd  $\wedge$  toBeC  $\neq$   $\emptyset$  do
    choose  $x \in$  toBeC;
    while true do
121:      $m := LL(\text{indir}[x]);$ 
122:      $\text{node}[mp] := \text{node}[m];$ 
123:     if round[mp] = rnd + 1  $\wedge$  color[mp] = grey then
124:         color[mp] := white;
125:         if SC(indir[x], mp) then toBeC := toBeC - {x}; mp := m; break; fi;
126:         elseif VL(indir[x]) then toBeC := toBeC - {x}; break; fi;
    od;
    od;
127: CAS(shRnd, rnd, rnd + 1);
128: return;
end GCollect.

proc Mark_stack( $x: 1 \dots N$ ) =
    local  $w, y: 1 \dots N$ ;  $\text{suc}: \text{Bool}$ ;  $j, k: \mathbb{N}$ ;
    stack: Stack; head:  $\mathbb{N}$ ; set: a subset of  $1 \dots N$ ;
    ch:  $[1 \dots C]$  of  $1 \dots N$ ;  $m, n: 1 \dots N + P$ ;
150: toBeC := toBeC - {x}; set := {x}; head := 0;
151: while shRnd = rnd  $\wedge$  set  $\neq$   $\emptyset$  do
    choose  $w \in$  set;
    while true do
152:      $m := LL(\text{indir}[w]);$ 
153:      $\text{node}[mp] := \text{node}[m];$ 
154:     if color[mp] = grey  $\wedge$  round[mp] = rnd + 1 then
155:          $k := \text{ari}[mp];$ 
        for  $j := 1$  to  $k$  do ch[j] := child[mp, j]; od;
156:     if VL(indir[w]) then
        set := set - {w}; head++; stack[head] := w; j := 1;
157:     while shRnd = rnd  $\wedge$   $j \leq k$  do
         $y := \text{ch}[j];$ 
        if  $y \in$  toBeC then
            while true do
158:                  $n := LL(\text{indir}[y]);$ 
159:                  $\text{node}[mp] := \text{node}[n];$ 
160:                 if (father[mp] = -1  $\vee$  father[mp] = w)
                     $\wedge$  round[mp] = rnd + 1 then
161:                     if father[mp] = -1 then father[mp] := w; fi;
162:                     if SC(indir[y], mp) then
                        toBeC := toBeC - {y}; mp := n;
                        set := set + {y}; break; fi;
                    elseif VL(indir[y]) then break; fi;
                od; fi;
             $j := j + 1;$ 
        od;
        break; fi;
    elseif VL(indir[w]) then set := set - {w}; break; fi;
    od;
od;
170: while shRnd = rnd  $\wedge$  head  $\neq$  0 do
     $y := \text{stack}[\text{head}];$ 
    while true do
171:      $m := LL(\text{indir}[y]);$ 

```

```

172:     node[mp] := node[m];
173:     if color[mp] = grey ∧ round[mp] = rnd + 1 then
174:         color[mp] := black;
           srcnt[mp] := srcnt[mp] - freecnt[mp]; freecnt[mp] := 0;
175:         if SC(indir[y], mp) then mp := m; head--; break; fi;
176:         elseif VL(indir[y]) then head--; break; fi;
           od;
       od;
180: return;
end Mark_stack.

```

```

proc Create(): 1...N =
  local m: 1...N + P; x: 1...N;
  while true do
200:   choose x ∈ 1...N;
201:   m := LL(indir[x]);
202:   node[mp] = node[m];
203:   if color[mp] = white then
204:     color[mp] := black; srcnt[mp] := 1; arity[mp] := 0;
205:     if SC(indir[x], mp) then
           roots := roots + {x};
           mp := m; break; fi;
206:     elseif time to do GC then
           GCollect(); fi;
       od;
207: return x
end Create.

```

```

proc AddChild(x, y: 1...N): Bool =
  {R(self, x) ∧ R(self, y)}
  local m: 1...N + P; suc: Bool;
250: suc := false;
  while true do
251:   m := LL(indir[x]);
252:   node[mp] := node[m];
253:   if arity[mp] < C then
254:     arity[mp]++;
           child[mp, arity[mp]] := y;
255:     if SC(indir[x], mp) then
           mp := m; suc := true; break; fi;
256:     elseif VL(indir[x]) then break; fi;
       od;
257: return suc
end AddChild.

```

```

proc GetChild(x: 1...N, rth: 1...N): 0...N =
  {R(self, x)}
  local m: 1...N + P; y: 1...N;
  while true do
280:   m := LL(indir[x]);
281:   node[mp] := node[m];
282:   if 1 ≤ rth ≤ arity[mp] then y := child[mp, rth]; else y := 0; fi;
283:   if VL(indir[x]) then break; fi;
       od;
284: return y

```

end *GetChild*.

```

proc Make(c: array [1...C] of 1...N, n: 1...C): 1...N =
{ $\forall j: 1 \dots n: R(\text{self}, c[j])$ }
  local m: 1...N + P; x: 1...N; j:  $\mathbb{N}$ ;
  while true do
300:   choose x  $\in$  [1...N];
301:   m := LL(indir(x));
302:   node[mp] := node[m];
303:   if color[mp] = white then
304:     color[mp] := black;
       srcnt[mp] := 1; arity[mp] := n;
       for j := 1 to n do child[mp, j] := c[j] od;
305:   if SC(indir(x), mp) then
       roots := roots + {x};
       mp := m; break; fi;
306:   elseif time to do GC then
       GCollect(); fi;
  od;
307: return x
end Make.

```

```

proc Protect(x: 1...N) =
{ $R(\text{self}, x) \wedge x \notin \text{roots}$ }
  local m: 1...N + P;
  while true do
400:   m := LL(indir[x]);
401:   node[mp] := node[m];
402:   srcnt[mp]++;
403:   if SC(indir[x], mp) then
       roots := roots + {x};
       mp := m; break; fi;
  od;
404: return
end Protect.

```

```

proc UnProtect(z: 1...N) =
{ $z \in \text{roots}$ }
  local m: 1...N + P;
  while true do
450:   m := LL(indir[z]);
451:   node[mp] := node[m];
452:   freecnt[mp]++;
453:   if SC(indir[x], mp) then
       roots := roots \ {z};
       mp := m; break; fi;
  od;
454: return
end UnProtect.

```

```

proc Send(x: 1...N, r: 1...P) =
{ $R(\text{self}, x) \wedge \text{Mbox}[\text{self}, r] = 0$ }
  local m: 1...N + P;
  while true do
500:   m := LL(indir[x]);

```

```

501:   node[mp] := node[m];
502:   srcnt[mp]++;
503:   if SC(indir[x], mp) then
504:       mp := m;
504:       Mbox[self, r] := x; break; fi;
       od;
505:   return
end Send.

```

```

proc Receive(r: 1...P): 0...N =
{Mbox[r, self] ≠ 0}
  local x: 1...N;
550:  x := Mbox[r, self];
551:  if x ∉ roots then
      roots := roots ∪ {x};
      Mbox[r, self] := 0;
  else
      while true do
552:         m := LL(indir[x]);
553:         node[mp] := node[m];
554:         srcnt[mp]--;
555:         if SC(indir[x], mp) then
            mp := m;
556:         Mbox[r, self] := 0;
            break; fi;
          od; fi;
557:  return
end Receive.

```

```

proc Check(r, q: 1...P): Bool
  local suc: Bool;
600:  suc := (Mbox[r, q] = 0);
601:  return suc;
end Receive.

```

References

- [1] H. Azatchi, Y. Levanoni, H. Paz, and E. Petrank. An on-the-fly mark and sweep garbage collector based on sliding views. In *Proceedings of the 18th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 269–281. ACM Press, 2003.
- [2] R.J.R. Back and J. von Wright. Stepwise refinement of distributed systems: Models, formalism, correctness: Refinement calculus. In J.W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Stepwise Refinement of Distributed Systems*, volume 430 of *Lecture Notes in Computer Science*, pages 42–93. Springer-Verlag, 1990.
- [3] G. Barnes. A method for implementing lock-free data structures. In *Proceedings of the 5th ACM Symposium on Parallel Algorithms and Architectures*, pages 261–270, June 1993.
- [4] M. Ben-Ari. Algorithms for on-the-fly garbage collection. *ACM Transactions on programming Languages and Systems*, 6(3):333–344, 1984.
- [5] B.N. Bershad. Practical considerations for non-blocking concurrent objects. In *Proceedings of the Thirteenth International Conference on Distributed Computing Systems*, pages 264–274, 1993.

- [6] H. Boehm, A. J. Demers, and S. Shenker. Mostly parallel garbage collection. In *Proceedings of the ACM SIGPLAN 1991 conference on Programming language design and implementation*, pages 157–164. ACM Press, 1991.
- [7] K. M. Chandy. *Parallel program design: a foundation*. Addison-Wesley Longman Publishing Co., Inc., 1988.
- [8] E. Cohen and L. Lamport. Reduction in TLA. In *Proceedings of the 9th International Conference on Concurrency Theory*, pages 317–331, 1998.
- [9] C. Cornes, J. Courant, and et al. The coq proof assistant - reference manual v 6.1, 1997.
- [10] D.L. Detlefs, P.A. Martin, M. Moir, and G.L. Steele Jr. Lock-free reference counting. *Distributed Computing*, 15(4):255–71, December 2002.
- [11] 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. *Communications of the ACM*, 21(11):966–975, November 1978.
- [12] D. Doligez and X. Leroy. A concurrent generational garbage collector for a multi-threaded implementation of ml. In *Proceedings of the 1993 ACM Symposium on Principles of Programming Languages*, pages 113–123, January 1993.
- [13] T. Endo, K. Taura, and A. Yonezawa. A scalable mark-sweep garbage collector on large-scale shared-memory machines. In *Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)*, pages 1–14. ACM Press, 1997.
- [14] C. Flood, D. Detlefs, N. Shavit, and C. Zhang. Parallel garbage collection for shared memory multiprocessors. In *Usenix Java Virtual Machine Research and Technology Symposium (JVM '01)*, Monterey, CA, April 2001.
- [15] H. Gao, J.F. Groote, and W.H. Hesselink. Efficient almost wait-free parallel accessible dynamic hashtables. Technical Report CS-Report 03-03, Eindhoven University of Technology, The Netherlands, 2003.
- [16] H. Gao, J.F. Groote, and W.H. Hesselink. Almost wait-free resizable hashtables (extended abstract). In *Proceedings of 18th International Parallel & Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, April 2004.
- [17] H. Gao and W.H. Hesselink. A formal reduction for lock-free parallel algorithms. In *Proceedings of the 16th Conference on Computer Aided Verification (CAV)*, July 2004.
- [18] H. Gao and W.H. Hesselink. A general lock-free algorithm using compare-and-swap. http://www.cs.rug.nl/~wim/mechver/lockfree_reduction, October 2004.
- [19] J.F. Groote, W.H. Hesselink, S. Mauw, and R. Vermeulen. An algorithm for the asynchronous write-all problem based on process collision. *Distributed Computing*, 14:75–81, 2001.
- [20] K. Havelund. Mechanical verification of a garbage collector. In José Rolim et al., editors, *Parallel and Distributed Processing (Combined Proceedings of 11 Workshops)*, volume 1586 of *Lecture Notes in Computer Science*, pages 1258–1283. Springer-Verlag, April 1999. Presented at the Workshop on Formal Methods for Parallel Programming: Theory and Applications (FMPPTA).
- [21] M. Herlihy. A methodology for implementing highly concurrent data objects. *ACM Transactions on Programming Languages and Systems*, 15(5):745–770, November 1993.

- [22] M.P. Herlihy, V. Luchangco, and M. Moir. The repeat offender problem: A mechanism for supporting dynamic-sized, lock-free data structure. In *Proceedings of 16th International Symposium on Distributed Computing*, pages 339–353. Springer-Verlag, October 2002.
- [23] M.P. Herlihy and J.E.B. Moss. Lock-free garbage collection for multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 3(3):304–311, 1992.
- [24] W.H. Hesselink. http://www.cs.rug.nl/~wim/mechver/garbage_collection
- [25] W.H. Hesselink. Using eternity variables to specify and prove a serializable database interface. *Science of Computer Programming*, 51:47–85, 2004.
- [26] W.H. Hesselink and J.F. Groote. Wait-free concurrent memory management by Create, and Read until Deletion. *Distributed Computing*, 14(1):31–39, January 2001.
- [27] R. L. Hudson and J. E. B. Moss. Sapphire: copying gc without stopping the world. In *ISCOPE Conference on ACM 2001 Java Grande*, pages 48–57. ACM Press, 2001.
- [28] L. Huelsbergen and J. R. Larus. A concurrent copying garbage collector for languages that distinguish (im)mutable data. In *Proceedings of the fourth ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 73–82. ACM Press, 1993.
- [29] E.H. Jensen, G.W. Hagensen, and J.M. Broughton. A new approach to exclusive data access in shared memory multiprocessors. Technical Report UCRL-97663, Lawrence Livermore National Laboratory, January 1987.
- [30] R. Jones and R. Lins. *Garbage collection: algorithms for automatic dynamic memory management*. John Wiley & Sons, Inc., 1996.
- [31] Richard Jones. *Garbage Collection: Algorithms for Automatic Dynamic Memory Management*. John Wiley and Sons, July 1996. With a chapter on Distributed Garbage Collection by Rafael Lins. Reprinted 1997 (twice), 1999, 2000.
- [32] J.E. Jonker. On-the-fly garbage collection for several mutators. *Distributed Computing*, 5:187–199, 1992.
- [33] P.C. Kanellakis and A. A. Shvartsman. *Fault-Tolerant Parallel Computation*. Kluwer Academic Publishers, 1997.
- [34] L. Lamport. The temporal logic of actions. *ACM Transactions on Programming Languages and Systems*, 16(3):872–923, 1994.
- [35] Y. Levanoni and E. Petrank. An on-the-fly reference counting garbage collector for java. In *Proceedings of the 16th ACM SIGPLAN conference on Object oriented programming, systems, languages, and applications*, pages 367–380. ACM Press, 2001.
- [36] V. Luchangco, M. Moir, and N. Shavit. Nonblocking k-compare-single-swap. In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pages 314–323. ACM Press, 2003.
- [37] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, 1996.
- [38] H. Massalin and C. Pu. A lock-free multiprocessor os kernel. Technical Report CUCS-005-91, Columbia University, 1991.
- [39] M. M. Michael. Safe memory reclamation for dynamic lock-free objects using atomic reads and writes. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 21–30. ACM Press, 2002.

- [40] M.M. Michael. High performance dynamic lock-free hash tables and list-based sets. In *Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures*, pages 73–82. ACM Press, 2002.
- [41] M. Moir. Practical implementations of non-blocking synchronization primitives. In *Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing*, pages 219–228. ACM Press, 1997.
- [42] A.J. Mooij. Non-blocking implementations of LL, VL and SC. Private Communication, 2004.
- [43] L. Moreau and J. Duprat. A construction of distributed reference counting. *Acta Inf.*, 37(8):563–595, 2001.
- [44] J. O’Toole and S. Nettles. Concurrent replicating garbage collection. In *Proceedings of the 1994 ACM conference on LISP and functional programming*, pages 34–42. ACM Press, 1994.
- [45] S. Owicki and L. Lamport. Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.*, 4(3):455–495, 1982.
- [46] S. Owre, N. Shankar, J.M. Rushby, and D.W.J. Stringer-Calvert. *PVS Version 2.4: System Guide, Prover Guide, PVS Language Reference*, 2001.
- [47] H. Rodrigues and R. Jones. Cyclic distributed garbage collection with group merger. In *Proceedings of 12th European Conference on Object-Oriented Programming, ECOOP98*, pages 249–273, Brussels, July 1998. Springer.
- [48] H. Sundell and P. Tsigas. Scalable and lock-free concurrent dictionaries. In *Proceedings of the 2004 ACM Symposium on Applied computing*, pages 1438–1445, 2004.
- [49] J.D. Valois. Lock-free linked lists using compare-and-swap. In *Proceedings of the fourteenth annual ACM symposium on Principles of distributed computing*, pages 214–222. ACM Press, 1995.