

Long-lived Rambo: Trading knowledge for communication[☆]

Chryssis Georgiou^a, Peter M. Musial^b, Alexander A. Shvartsman^{b,c,*}

^a *Department of Computer Science, University of Cyprus, 1678 Nicosia, Cyprus*

^b *Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06278, USA*

^c *Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

Abstract

Shareable data services providing consistency guarantees, such as atomicity (linearizability), make building distributed systems easier. However, combining linearizability with efficiency in practical algorithms is difficult. A reconfigurable linearizable data service, called RAMBO, was developed by Lynch and Shvartsman. This service guarantees consistency under dynamic conditions involving asynchrony, message loss, node crashes, and new node arrivals. The specification of the original algorithm is given at an abstract level aimed at concise presentation and formal reasoning about correctness. The algorithm propagates information by means of gossip messages. If the service is in use for a long time, the size and the number of gossip messages may grow without bound. This paper presents a consistent data service for *long-lived* objects that improves on RAMBO in two ways: it includes an incremental communication protocol and a leave service. The new protocol takes advantage of the local knowledge, and carefully manages the size of messages by removing redundant information, while the leave service allows the nodes to leave the system gracefully. The new algorithm is formally proved correct by forward simulation using levels of abstraction. An experimental implementation of the system was developed for networks-of-workstations. The paper also includes selected analytical and preliminary empirical results that illustrate the advantages of the new algorithm.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Distributed algorithms; Atomic memory service; Long-lived executions

1. Introduction

This paper presents a practical algorithm implementing long-lived, survivable, atomic read/write objects in dynamic networks, where participants may join, leave, or fail during the course of computation. The survivability of data is ensured through redundancy: the data are replicated and maintained at several network locations. Replication introduces the challenges of maintaining *consistency* among the replicas, and managing *dynamic participation* as the collections of network locations storing the replicas change due to arrivals, departures, and failures of nodes.

A new approach to implementing atomic read/write objects for dynamic networks was developed by Lynch and Shvartsman [1] and extended by Gilbert et al. [2]. They developed a memory service called RAMBO (Reconfigurable Atomic Memory for Basic Objects) that maintains atomic, a.k.a. linearizable, readable/writable data in highly dynamic

[☆] This work is supported in part by the NSF Grants 9988304, 0121277, and 0311368, and by the NSF CAREER Award 9984778.

* Corresponding author at: Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06278, USA.
E-mail address: aas@cse.uconn.edu (A.A. Shvartsman).

environments. In order to achieve availability in the presence of failures, the objects are replicated at several network locations. In order to maintain consistency in the presence of small and transient changes, the algorithm uses *configurations* of locations, each of which consists of a set of *members* plus sets of *read-quorums* and *write-quorums*. In order to accommodate larger and more permanent changes, the algorithm supports *reconfiguration*, by which the set of members and the sets of quorums are modified. Any configuration may be installed at any time. Obsolete configurations can be removed from the system without interfering with the ongoing read and write operations. The algorithm tolerates arbitrary patterns of asynchrony, node crashes, and message loss. It is formally shown [2,1] that atomicity is maintained in any execution of the algorithm. We developed an experimental implementation of this algorithm on a network-of-workstations [3].

The original RAMBO algorithm is formulated at an abstract level aimed at concise specification and formal reasoning about the algorithm's correctness. Consequently, the algorithm incorporates a simple communication protocol that maintains very little in the way of protocol states. Nevertheless, showing correctness requires careful arguments involving subtle race conditions [2,1]. The algorithm propagates information among the participants by means of gossip messages that contain information representing the sender's state. The number and the size of gossip messages may in fact grow without bound. This renders the algorithm impractical for use in *long-lived* applications.

The gossip messages in RAMBO include the set of participants, and the size of these messages increases over time for two reasons. First, RAMBO allows new participants to join the computation, but it does not allow the participants to leave gracefully. In order to leave, the participants must pretend to crash. Given that in asynchronous systems failure detection is difficult, it may be impossible to distinguish departed nodes from the nodes that crash. Second, RAMBO gossips information among the participants without regard for what may already be known at the destination. Thus, a participant will repeatedly gossip substantial amounts of information to others even if it has not learned anything new since the last time it gossiped. While such redundant gossip helps with tolerating message loss, it substantially increases the communication burden. Given that the ultimate goal for this algorithm is to be used in long-lived applications, and in dynamic networks with an unknown (and possibly infinite) universe of nodes, the algorithm must be carefully refined to substantially improve its communication efficiency.

Contributions. The paper presents a new algorithm for reconfigurable atomic memory for dynamic networks. The algorithm, called LL-RAMBO, makes implementing atomic survivable objects practical in long-lived systems by managing the knowledge accumulated by the participants and the size of the gossip messages. Each participating node maintains a more complicated protocol state and, with the help of additional local processing, this investment is traded for substantial reductions in the size and the number of gossip messages. Based on [2,1], we use Input/Output Automata (IOA) [4] to specify the algorithm, and then prove it correct in two stages by forward simulation, using levels of abstraction. We include analytical and preliminary empirical results illustrating the advantages of the new algorithm. In more detail, our contributions are as follows.

- (1) We develop L-RAMBO that implements an atomic memory service and includes a *leave* service (Section 4). We prove the correctness (safety) of L-RAMBO by forward simulation of RAMBO; hence we show that every trace of L-RAMBO is a trace of RAMBO.
- (2) We develop LL-RAMBO by refining L-RAMBO to implement *incremental gossip* (Section 5). We prove that LL-RAMBO implements the atomic service by forward simulation of L-RAMBO. This shows that every trace of LL-RAMBO is a trace of L-RAMBO, and thus a trace of RAMBO. The proof involves subtle arguments relating the knowledge extracted from the local state to the information that is *not* included in gossip messages. We present the proof in two steps for two reasons: (i) the presentation matches the intuition that the leave service and the incremental gossiping are independent; and (ii) the resulting proof is simpler than a direct simulation of RAMBO by LL-RAMBO.
- (3) We show (Section 6) that LL-RAMBO consumes fewer communication resources than RAMBO, while preserving the same read and write operation latency, which under certain steady-state assumptions is at most $8d$ time, where d is the maximum message delay unknown to the algorithm. Under these assumptions, in runs with periodic gossip, LL-RAMBO achieves substantial reductions in communication.
- (4) We implemented all algorithms on a network-of-workstations. Preliminary empirical results complement the analytical comparison of the two algorithms (Section 7).

Background. Several approaches can be used to implement consistent data in (static) distributed systems. Starting with Gifford [5] and Thomas [6], many algorithms used collections of intersecting sets of object replicas to solve the

consistency problem. Upfal and Wigderson [7] use majority sets of readers and writers to emulate shared memory. Awerbuch and Vitanyi [8] use matrices of registers where the rows and the columns are written and respectively read by specific processors. Attiya, Bar-Noy and Dolev [9] use majorities to implement shared objects in static message passing systems. Extensions with reconfigurable quorums have been explored [10,11]. These systems have limited ability to support long-lived data when the longevity of processors is limited. Virtual synchrony [12], and group communication services (GCS) in general [13], can be used to implement consistent objects, e.g., by using a global totally ordered broadcast. The set of nodes participating in a GCS can evolve; however, forming a new view is initiated after a single failure and can take a substantial time, while reads and writes are delayed during view formation. In our algorithm, as in [1,2], reads and writes can make progress during reconfiguration. In the current approach, arbitrary new configurations can be introduced. This yields a more dynamic system compared to [14–18] that would require that some new quorums include nodes from the old quorums, thus restricting the choice of the new configuration through the static constraints that need to be satisfied even before the reconfiguration.

The work on reconfigurable atomic memory [10,1,2] results in algorithms that are more dynamic because they place fewer restrictions on the choice of new configurations and allow for the universe of processors to evolve arbitrarily. However, these approaches are based on abstract communication protocols that are not suited for long-lived systems. In this paper, we provide a long-lived solution by introducing graceful processor departures and incremental gossip. The idea of incrementally propagating information among participating nodes has been previously used in a variety of different settings, e.g., [19–25]. Incremental gossip is also called anti-entropy [26,27] or reconciliation [28]; these concepts are used in database replication algorithms; however, due to the nature of the application they assume stronger assumptions, e.g., ordering of messages.

Document structure. The model and the data types used throughout the paper are introduced in Section 2. In Section 3, we review RAMBO. In Section 4 we specify and prove correct the graceful leave service. Section 5 presents our final solution, a system with leave and incremental gossip, and proves it correct. In Section 6 we present the analytical performance analysis, and in Section 7 we demonstrate the experimental results. We conclude in Section 8. A preliminary version of this paper appears in [29].

2. Model and definitions

We now define the notion of memory atomicity and the model of computation used in this paper. We also define common data types and quorum systems.

2.1. Atomicity

We assume that the clients of the data service for each object x are well-formed, in that each client invokes only one read or write operation on the object at a time.

The data object x is atomic (a.k.a. linearizable) if in any execution, whenever all the read and write operations that are invoked complete, then the read and write operations for object x can be partially ordered by an ordering $<$, so that the following conditions are satisfied:

- (1) The partial order is consistent with the external order of invocations and responses; that is, there do not exist read or write operations π_1 and π_2 such that π_1 completes before π_2 starts, yet $\pi_2 < \pi_1$.
- (2) All write operations are totally ordered and every read operation is ordered with respect to all the writes.
- (3) Every read operation ordered after any write returns the value of the last write preceding it in the partial order; any read operation ordered before all writes returns $(v_0)_x$.

(Atomicity is often defined in terms of an equivalence with a serial memory. The definition given here implies this equivalence, as shown, for example, in Lemma 13.16 in [30].)

2.2. Model of computation

We assume the asynchronous distributed, message-passing model with an unknown number of processors that have unique identifiers. Processors may join the service at any time after the service is initiated, and may crash at any time during the computation, and can do so without warning.

The non-failed processors communicate via point-to-point channels with the following properties: (i) messages are not corrupted, (ii) messages may be lost, (iii) messages may be delivered in an arbitrary order, and (iv) messages are not spontaneously generated by the channels, which means that for each receive event there is a corresponding send event. We denote by $Channel_{i,j}$ the channel from node i to node j .

In the correctness (safety) analysis, we assume complete asynchrony (Section 4). For the purpose of performance analysis (Section 6) we impose bounds on the duration of processing steps and the delivery of messages. These bounds are unknown to the algorithm.

2.3. Data types

In the discussion that follows, we assume two distinguished elements, \perp and \pm , which are not in any of the basic types. For any type A , we define new types $A_{\perp} = A \cup \{\perp\}$ and $A_{\pm} = A \cup \{\perp, \pm\}$. If A is a partially ordered set, we augment its ordering by assuming that $\perp < a < \pm$ for every $a \in A$.

We define C to be the set of *configuration identifiers* and assume only the trivial partial order on C , in which all elements are incomparable. For each $c \in C$ we define: (i) *members*(c), a finite subset of node identifiers, (ii) *read-quorums*(c), a set of finite subsets of *members*, and (iii) *write-quorums*(c), a set of finite subsets of *members*. We require that for every $R \in \text{read-quorums}(c)$, and every $W \in \text{write-quorums}(c)$, $R \cap W \neq \emptyset$. No intersection requirement is imposed on the sets of members or on the quorums from distinct configurations.

In Table 1 we define data types, distinguished elements, and functions used in the sequel. The definitions of variables and data types are also recalled throughout the discussion to provide intuition into how they are used.

3. Reconfigurable atomic memory for basic objects (RAMBO)

We now describe the RAMBO algorithm as presented in [1], including the rapid configuration upgrade as given in [2]. The algorithm is given for a single object. Atomicity is preserved under composition, and multiple objects are composed to yield a complete shared memory. Therefore, throughout the rest of the paper, we describe an implementation for a particular object x and suppress any explicit mention of x . Thus, we write V , v_0 , c_0 , and i_0 from now on as shorthand for V_x , $(v_0)_x$, $(c_0)_x$, and $(i_0)_x$, respectively. For the detailed Input/Output Automata code of the algorithms see Appendix and [1,2].

In order to achieve fault tolerance and availability, RAMBO replicates objects at several network locations. In order to maintain memory consistency in the presence of small and transient changes, the algorithm uses *configurations*. In order to accommodate larger and more permanent changes, the algorithm supports *reconfiguration*, by which the set of members and the sets of quorums are modified. Any quorum configuration may be installed, and atomicity is preserved in all executions.

The algorithm consists of three kinds of automata: (i) *Joiner* automata, handling join requests; (ii) *Recon* automata, handling reconfiguration requests (*recon*), and generating a totally ordered sequence of configurations; and (iii) *Reader–Writer* automata, handling *read* and *write* requests, manage configuration upgrades, and implement gossip messaging. The overall system is the composition of these automata with the automata modelling point-to-point communication channels, see Fig. 1. (Properties of the channel automata are described in Section 2.)

The *Joiner* automaton is quite straightforward, simply sending a *join* message when node i joins, and sending a *join-ack* message whenever a *join* message is received. The *Recon* automaton establishes a total ordering of configurations; it is implemented using a consensus service (e.g., Paxos [31]), or it can be a much simpler service when the set of possible configurations is finite [32]. Here, we assume that a total ordering exists, and we need not discuss this further (for details see [1]).

The external signature of the service is in Fig. 2. A client at node i uses a join_i action to join the system. After receiving join-ack_i , the client can issue read_i and write_i requests, which result in read-ack_i and write-ack_i responses. The client can issue a recon_i request a reconfiguration. The fail_i action models a crash at node i .

Remark. The fail_i action is an input action coming from the environment. In the original RAMBO, it is used solely to signal the crash of node i . In this paper, we introduce a new kind of action coming from the environment—the leave_i action. This is discussed in detail in the next section. In showing the correctness of the new algorithm, we assume that fail_i and leave_i are synonymous: both come from the environment and both result in node i ceasing to participate in the computation. Thus, we model graceful departures as a form of failure.

Table 1
Definitions of data types

I	The totally-ordered set of <i>locations</i> , i.e., node identifiers.
T	The set of <i>tags</i> , defined as $\mathbb{N} \times I$.
M	The set of <i>messages</i> , which is a set of all legal messages that may be sent by the algorithm.
X	The set of <i>object identifiers</i> , partitioned into subsets $X_i, i \in I$. X_i is the set of identifiers for objects that may be created at location i . For any $x \in X$, $(i_0)_x$ denotes the unique i such that $x \in X_i$.
V_x	For each $x \in X$: (i) V_x , the set of values that object x may take on, and (ii) $(v_0)_x \in V_x$, the initial value of x .
C	The set of <i>configuration identifiers</i> . For each $c \in C$ we define: (i) <i>members</i> (c), a finite subset of I . (ii) <i>read-quorums</i> (c), a set of finite subsets of <i>members</i> (c). (iii) <i>write-quorums</i> (c), a set of finite subsets of <i>members</i> (c). For every c , every $R \in \text{read-quorums}(c)$, and every $W \in \text{write-quorums}(c)$, $R \cap W \neq \emptyset$.
$(c_0)_x$	For each $x \in X$, $(c_0)_x \in C$, the <i>initial configuration identifier</i> for x . We assume that <i>members</i> ($(c_0)_x$) = $\{(i_0)_x\}$. That is, the initial configuration for object x has only a single member, which is the creator of x .
<i>update</i>	A binary function on C_{\pm} , defined by $\text{update}(c, c') = \max(c, c')$ if c and c' are comparable (in the augmented partial ordering of C_{\pm}), $\text{update}(c, c') = c$ otherwise.
<i>extend</i>	A binary function on C_{\pm} , defined by $\text{extend}(c, c') = c'$ if $c = \perp$ and $c' \in C$, and $\text{extend}(c, c') = c$ otherwise.
$CMap$	The set of <i>configuration maps</i> , defined as the set of mappings from \mathbb{N} to C_{\pm} , $\mathbb{N} \rightarrow C_{\pm}$. We extend the <i>update</i> and <i>extend</i> operators elementwise to binary operations on $CMap$.
<i>truncate</i>	A unary function on $CMap$, defined by $\text{truncate}(cm)(k) = \perp$ if there exists $\ell \leq k$ such that $cm(\ell) = \perp$, $\text{truncate}(cm)(k) = cm(k)$ otherwise. This truncates configuration map cm by removing all the configuration identifiers that follow a \perp .
<i>Truncated</i>	The subset of $CMap$ such that $cm \in \text{Truncated}$ if and only if $\text{truncate}(cm) = cm$.
<i>Usable</i>	The subset of $CMap$ such that $cm \in \text{Usable}$ iff the pattern occurring in cm consists of a prefix of finitely many $\pm s$, followed by an element of C , followed by an infinite sequence of elements of C_{\perp} in which all but finitely many elements are \perp .

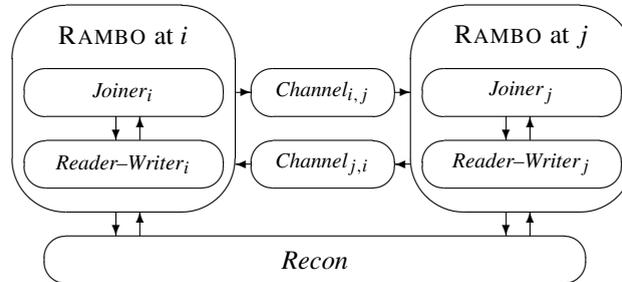


Fig. 1. RAMBO architecture depicting automata at nodes i and j , the channels, and the *Recon* service.

Input:

join(rambo, J) _{i} , J a finite subset of $I - \{i\}$,
 $i \in I$, such that if $i = i_0$ then $J = \emptyset$
 read _{i} , $i \in I$
 write(v) _{i} , $v \in V, i \in I$
 recon(c, c') _{i} , $c, c' \in C, i \in \text{members}(c), i \in I$
 fail _{i} , $i \in I$

Output:

join-ack(rambo) _{i} , $i \in I$
 read-ack(v) _{i} , $v \in V, i \in I$
 write-ack _{i} , $i \in I$
 recon-ack(b) _{i} , $b \in \{ok, nok\}, i \in I$
 report(c) _{i} , $c \in C, i \in I$

Fig. 2. RAMBO: External signature.

Every node of the system maintains a *tag* and a *value* for the data object. Every time a new value is written, it is assigned a unique tag, with ties broken by process-ids. These tags are used to determine an ordering of the write operations, and therefore determine the value that a read operation should return. Read and write operations have

Table 2
State variables of *Reader-Writer_i*

<i>status</i>	A member of the set {idle, joining, active}, initially idle. Indicates status of the service at the local node.
<i>world</i>	A finite subset of I , initially \emptyset . The totally-ordered set of <i>locations</i> , i.e. node identifiers.
<i>value</i>	A member of V , initially v_0 , otherwise the value of the replica that was returned during the most recent read or write operation.
<i>tag</i>	A member of T , initially $(0, i_0)$, otherwise the value of the last tag that was returned during the most recent read or write operation.
<i>cmap</i>	A member of $CMap$, initially $cmap(0) = c_0$, $cmap(k) = \perp$ for $k \geq 1$.
<i>pnum1</i>	A natural number, \mathbb{N} , initially 0. Keeps count of the locally initiated phases.
<i>pnum2</i>	A mapping $I \rightarrow \mathbb{N}$, initially everywhere 0. Keeps track of phase counts of other nodes in the system.
<i>failed</i>	A Boolean, initially false. If set to true, it signals that the node stops participation in the service.
<i>op</i>	A record with fields: (i) <i>type</i> \in {read, write} (ii) <i>phase</i> \in {idle, query, prop, done}, initially idle, (iii) <i>pnum</i> \in \mathbb{N} , (iv) <i>cmap</i> \in $CMap$, (v) <i>acc</i> , a finite subset of I , and (vi) <i>value</i> \in V . This record is used to track the progress of the current read/write operation.
<i>upg</i>	A record with fields: (i) <i>phase</i> \in {idle, query, prop}, initially idle, (ii) <i>pnum</i> \in \mathbb{N} , (iii) <i>cmap</i> \in $CMap$, (iv) <i>acc</i> , a finite subset of I , and (v) <i>target</i> \in \mathbb{N} . This record is used to track the progress of the current configuration upgrade operation.

two phases, *query* and *propagation*, each accessing certain quorums of replicas. Assume the operation is initiated at node i . First, in the query phase, node i contacts read quorums to determine the most recent known tag and value. Then, in the propagation phase, node i contacts write quorums. If the operation is a read operation, the second phase propagates the largest discovered tag and its associated value. If the operation is a write operation, node i chooses a new tag, strictly larger than every tag discovered in the query phase. Node i then propagates the new tag and the new value to a write quorum. Note that every operation accesses both read and write quorums.

Configurations go through three stages: proposal, installation, and upgrade. First, a configuration is *proposed* by a `recon` event. Next, if the proposal is successful, the *Recon* service achieves consensus on the new configuration, and notifies participants with `decide` events. When every non-failed member of the prior configuration has been notified, the configuration is *installed*. The configuration is *upgraded* when every configuration with a smaller index has been removed. Upgrades are performed by the configuration upgrade operations. Each upgrade operation requires two phases, a query phase and a propagate phase. The first phase contacts a read-quorum and a write-quorum from the old configurations, and the second phase contacts a write-quorum from the new configuration. All three operations, *read*, *write*, and *configuration upgrade*, are implemented using gossip messages.

The *cmap* is a mapping from integer indices to configurations $\cup\{\perp, \pm\}$, initially mapping every integer to \perp . It records which configurations are active, which have not yet been created, indicated by \perp , and which have already been removed, indicated by \pm . The total ordering on configurations determined by *Recon* ensures that all nodes agree on which configuration is stored in each position in *cmap*. We define $c(k)$ to be the configuration associated with index k .

The record *op* is used to store information about the current phase of an ongoing read or write operation, while *upg* is used for information about an ongoing configuration upgrade operation. A node can process read and write operations concurrently with configuration upgrade operations. The *op.cmap* subfield records the configuration map associated with the operation. For read or write operations, this consists of the node's *cmap* when a phase begins, augmented by any new configurations discovered during the phase. A phase completes when the initiator has exchanged information with quorums from every valid configuration in *op.cmap*. The *pnum* subfield records the phase number when the phase begins, allowing the initiator to determine which responses correspond to the phase. The *acc* subfield accumulates the ids of the nodes from which quorums have responded during the current phase.

A short description of the variables introduced in this section can be found in Table 2, and of the data types in Table 1.

Remark. RAMBO uses monotonically increasing phase numbers to detect fresh responses. The new phase number is computed by incrementing the previous phase number. Here, we assume that new phase numbers are computed by adding an arbitrary positive integer to the previous phase number. This is sufficient for the correctness of RAMBO, while substantially simplifying the proof of correctness of our algorithm.

4. RAMBO with graceful leave

Here we augment RAMBO with a *leave service* allowing the participants to depart gracefully. We prove that the new algorithm, called L-RAMBO, implements atomic memory.

Nodes participating in RAMBO communicate by means of gossip messages containing the latest object value and bookkeeping information that includes the set of known participants. RAMBO allows participants to fail or leave without warning. Since, in asynchronous systems, it is difficult or impossible to distinguish slow or departed nodes from crashed nodes, RAMBO implements gossip to all known participants, regardless of their status. In highly dynamic systems, this leads to (i) the size of gossip messages growing without bounds; and (ii) the number of messages sent in each round of gossip increasing as new participants join the computation.

L-RAMBO allows graceful node departures by letting a node that wishes to leave the system to send notification messages to an arbitrary subset of known participants. When another node receives such a notification, it marks the sender as departed, and stops gossiping to that node. The remaining nodes propagate the information about the departed nodes to the other participants, eventually eliminating gossip to nodes that departed gracefully.

Specification of L-RAMBO. We interpret the fail_i event as synonymous with the leave_i event—both are inputs from the environment and both result in node i stopping participation in all operations. The difference between fail_i and leave_i is strictly internal: leave_i allows a node to leave gracefully. The well-formedness conditions of RAMBO and the specifications of Joiner_i and Recon remain unchanged. The introduction of the leave service affects only the specification of the Reader-Writer_i automata. These changes for L-RAMBO are given in Fig. 3, except for the segments of code annotated with a parenthesized asterisk ($(*)$) should be disregarded until our targeted long-lived algorithm LL-RAMBO is presented in Section 5 (we combine the two specifications to avoid unnecessary repetition and simplify presentation).

The signature of Reader-Writer_i automaton is extended with actions $\text{recv}(\text{leave})_{j,i}$ and $\text{send}(\text{leave})_{i,j}$ used to communicate the graceful departure status. The state of Reader-Writer_i is extended with new state variables: departed_i , the set of nodes that left the system, as known at node i , leave-world_i , the set of nodes that node i can inform of its own departure, once it decides to leave and sets leave-world_i to $\text{world}_i - \text{departed}_i$.

The key algorithmic changes involve the actions $\text{recv}(m)_{j,i}$ and $\text{send}(m)_{i,j}$. The messages in the original RAMBO algorithm include: W the *world* of the sender, v the object and its tag t , cm the *cmap*, pns the phase number of the sender, and pnr the phase number of the receiver that is known to the sender. The gossip message m in L-RAMBO also includes D , a new parameter, initialized to the departed set of the sender.

We now detail the leave protocol. Assume that nodes i and j participate in the service, and node i wishes to depart following the leave_i event, whose effects set the state variable failed_i to *true* in Joiner_i , Recon_i , and Reader-Writer_i . The leave_i action at Reader-Writer_i (see Fig. 3) also initializes the set leave-world_i to the identifiers found in world_i , less those found in departed_i . Now Reader-Writer_i is allowed to send one leave notification to any node in leave-world_i . This is done by the $\text{send}(\text{leave})_{i,j}$ action that arbitrarily chooses the destination j from leave-world_i . Note that node i may nondeterministically choose the original fail_i action (see Appendix or [1]), in which case no notification messages are sent (this is the “non-graceful” departure).

When Reader-Writer_i receives a leave notification from node j , it adds j to its departed_i set. Node i sends gossip messages to all nodes in the set $\text{world}_i - \text{departed}_i$, which includes information about j ’s departure. When Reader-Writer_i receives a gossip message that includes the set D , it updates its departed_i set accordingly.

A short description of variables introduced in this section can be found in Table 3, whereas short descriptions of other variables can be found in Table 2, and of the data types in Table 1.

Atomicity of L-RAMBO service. The L-RAMBO system is the composition of all Reader-Writer_i and Joiner_i automata, the Recon service, and $\text{Channel}_{i,j}$ automata for all $i, j \in I$. We show the atomicity of L-RAMBO by forward simulation that proves that any trace of L-RAMBO is also a trace of RAMBO, and thus L-RAMBO implements atomic objects. The proof uses history variables, annotated with the symbol $[h]$ in Fig. 3.

For each i , we define $h\text{-msg}_i$ to be the history variable that keeps track of all messages sent by Reader-Writer_i automata. Initially, $h\text{-msg}_i = \emptyset$ for all $i \in I$. Whenever a message m is sent by i to some node $j \in I$ via $\text{Channel}_{i,j}$, we let $h\text{-msg}_i \leftarrow h\text{-msg}_i \cup \{m, i, j\}$. We define $h\text{-MSG}$ to be $\bigcup_{i \in I} h\text{-msg}_i$. (The remaining history variables are used in reasoning about LL-RAMBO, in Section 5.)

The following lemma states that only “good” messages are sent.

Signature:

As in RAMBO, plus new actions:

Input : $\text{leave}_i, \text{recv}(\text{leave})_{j,i}$ Output : $\text{send}(\text{leave})_{i,j}$ **State:**

As in RAMBO, plus new states:

 leave-world , a finite subset of I , initially \emptyset departed , a finite subset of I , initially \emptyset **Transitions at i :**Input $\text{recv}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$

Effect:

if $\neg \text{failed} \wedge \text{status} \neq \text{idle}$ then $\text{status} \leftarrow \text{active}$ $\text{world} \leftarrow \text{world} \cup W$ $\text{departed} \leftarrow \text{departed} \cup D$ $[h]\text{hrecv-}W(j, i, pnr) \leftarrow W$ (*) $[h]\text{hrecv-}D(j, i, pnr) \leftarrow D$ (*) $\text{ig}(j).\text{w-known} \leftarrow \text{ig}(j).\text{w-known} \cup W$ (*) $\text{ig}(j).\text{w-unack} \leftarrow \text{ig}(j).\text{w-unack} - W$ (*) $\text{ig}(j).\text{d-known} \leftarrow \text{ig}(j).\text{d-known} \cup D$ (*) $\text{ig}(j).\text{d-unack} \leftarrow \text{ig}(j).\text{d-unack} - D$ (*) if $\text{pnr} > \text{ig}(j).\text{p-ack}$ then(*) $\text{ig}(j).\text{w-known} \leftarrow \text{ig}(j).\text{w-known} \cup \text{ig}(j).\text{w-unack}$ (*) $\text{ig}(j).\text{w-unack} \leftarrow \text{world} - \text{ig}(j).\text{w-known}$ (*) $\text{ig}(j).\text{d-known} \leftarrow \text{ig}(j).\text{d-known} \cup \text{ig}(j).\text{d-unack}$ (*) $\text{ig}(j).\text{d-unack} \leftarrow \text{departed} - \text{ig}(j).\text{d-known}$ (*) $\text{ig}(j).\text{p-ack} \leftarrow \text{pnum1}$ if $t > \text{tag}$ then $(\text{value}, \text{tag}) \leftarrow (v, t)$ $\text{cmap} \leftarrow \text{update}(\text{cmap}, \text{cm})$ $\text{pnum2}(j) \leftarrow \max(\text{pnum2}(j), \text{pns})$ if $\text{op.phase} \in \{\text{query}, \text{prop}\} \wedge \text{pnr} \geq \text{op.pnum}$ then $\text{op.cmap} \leftarrow \text{extend}(\text{op.cmap}, \text{truncate}(\text{cm}))$ if $\text{op.cmap} \in \text{Truncated}$ then $\text{op.acc} \leftarrow \text{op.acc} \cup \{j\}$

else

 $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\text{op.acc} \leftarrow \emptyset$ $\text{op.cmap} \leftarrow \text{truncate}(\text{cmap})$ if $\text{upg.phase} \in \{\text{query}, \text{prop}\} \wedge \text{pnr} \geq \text{upg.pnum}$ then $\text{upg.acc} \leftarrow \text{upg.acc} \cup \{j\}$ Input $\text{recv}(\text{leave})_{j,i}$

Effect:

if $\neg \text{failed} \wedge \text{status} = \text{active}$ then $\text{departed} \leftarrow \text{departed} \cup \{j\}$ (*) $\text{ig} \in \text{IGMap}$, initially $\forall k \in I$,(*) $\text{ig}(k).\text{w-known} = \emptyset, \text{ig}(k).\text{w-unack} = \emptyset$,(*) $\text{ig}(k).\text{d-known} = \emptyset, \text{ig}(k).\text{d-unack} = \emptyset$,(*) $\text{ig}(k).\text{p-ack} = 0$ Output $\text{send}(\langle W, D, v, t, cm, pns, pnr \rangle)_{i,j}$

Precondition:

 $\neg \text{failed}$ $\text{status} = \text{active}$ $j \in (\text{world} - \text{departed})$ $(W, D, v, t, cm, pns, pnr) =$ (*) $(\text{world}$ $\neg \text{ig}(j).\text{w-known}$,(*) departed $\neg \text{ig}(j).\text{d-known}$, $\text{value}, \text{tag}, \text{cmap}, \text{pnum1}, \text{pnum2}(j))$

Effect:

 $\text{pnum1} \leftarrow \text{pnum1} + 1$ $[h]\text{h-msg} \leftarrow \text{h-msg} \cup$ (*) $(\langle W, D, v, t, cm, pns, pnr \rangle, i, j)$ (*) $[h]\text{hsent-}W(i, j, pns) \leftarrow W$ (*) $[h]\text{hsent-}D(i, j, pns) \leftarrow D$ (*) $[h]\text{hs-world}(i, j, pns) \leftarrow \text{world}$ (*) $[h]\text{hs-departed}(i, j, pns) \leftarrow \text{departed}$ (*) $[h]\text{hs-w-known}(i, j, pns) \leftarrow \text{ig}(j).\text{w-known}$ (*) $[h]\text{hs-d-known}(i, j, pns) \leftarrow \text{ig}(j).\text{d-known}$ (*) $[h]\text{hs-w-unack}(i, j, pns) \leftarrow \text{ig}(j).\text{w-unack}$ (*) $[h]\text{hs-d-unack}(i, j, pns) \leftarrow \text{ig}(j).\text{d-unack}$ (*) $[h]\text{hs-p-ack}(i, j, pns) \leftarrow \text{ig}(j).\text{p-ack}$ Input leave_i

Effect:

if $\neg \text{failed}$ then $\text{failed} \leftarrow \text{true}$ $\text{departed} \leftarrow \text{departed} \cup \{i\}$ $\text{leave-world} \leftarrow \text{world} - \text{departed}$ Output $\text{send}(\text{leave})_{i,j}$

Precondition:

 $j \in \text{leave-world}$

Effect:

 $\text{leave-world} \leftarrow \text{leave-world} - \{j\}$ Fig. 3. Modification of Reader-Writer_i for L-RAMBO, and for LL-RAMBO (the lines annotated with (*)).

Table 3

State variables that are used in L-RAMBO, in addition to the ones defined in Table 2

leave-world	A finite subset of I , initially \emptyset ; otherwise, contains node identifiers to which the leave notification message may be sent to.
departed	A finite subset of I , initially \emptyset ; otherwise, it is a subset of node identifiers found in world .
h-msg	$\forall \langle m, \cdot, i \rangle \in \text{h-msg}_i : m \in M$. A history variable that keeps track of all messages sent by node i . A set of all sent messages is denoted as $\text{h-MSG} = \bigcup_{\forall \ell} \text{h-msg}_\ell$.

Lemma 4.1. *In any execution of L-RAMBO, if m is a message received by node i in a $\text{recv}(m)_{i,j}$ event, then $\langle m, j, i \rangle \in \text{h-MSG}$, and $m \in \{\langle W, D, v, t, cm, pns, pnr \rangle, \text{leave}, \text{join}\}$, where $\langle W, D, v, t, cm, pns, pnr \rangle \in I \times I \times V \times T \times \text{Cmap} \times \mathbb{N} \times \mathbb{N}$.*

Proof. By the definition of $\text{Channel}_{j,i}$, the messages are not corrupted, and for every receive there exists a preceding send event (messages are not manufactured by the channel). Therefore, m must have been sent by some node $j \in I$ during some earlier event of the execution. Hence $\langle m, j, i \rangle \in \text{h-msg}_j$; by definition, $\langle m, j, i \rangle \in \text{h-MSG}$. Messages are sent only in Reader-Writer_j automaton's $\text{send}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$ or $\text{send}(\text{leave})_{j,i}$ events, or

in $Joiner_j$ automaton's $\text{send}(\text{join})_{j,i}$ event, and by the code of these events $m = \langle W, D, v, t, cm, pns, pnr \rangle$, $m = \text{leave}$, or $m = \text{join}$. \square

Definition 4.2 and Theorem 4.3 are used to show the safety of L-RAMBO.

Definition 4.2 ([30]). Let A and S be two automata with the same external signature. A **forward simulation** from A to S is a relation $R \subseteq \text{states}(A) \times \text{states}(S)$ that satisfies the following two conditions:

- (1) If t is any initial state of A , then there is an initial state s of S such that $s \in R(t)$, where $R(t)$ is an abbreviation for $\{s: (t, s) \in R\}$.
- (2) If t and $s \in R(t)$ are reachable states of A and S respectively, and if (t, π, t') is a step of A , then there exists an execution fragment of S from s to some $s' \in R(t')$, having the same trace as step (t, π, t') .

Theorem 4.3 ([30]). *If there is a simulation from automaton A to automaton S , then $\text{traces}(A) \subseteq \text{traces}(S)$.*

Next, we show that L-RAMBO implements RAMBO, assuming the environment behaves as (informally) described in Section 3. Showing well-formedness is straightforward by inspecting the code. The proof of atomicity is based on a forward simulation from L-RAMBO to RAMBO.

Theorem 4.4. *L-RAMBO implements atomic read/write objects.*

Proof. We show that L-RAMBO algorithm simulates RAMBO. Specifically, we show that there exists a simulation relation R from L-RAMBO to RAMBO that satisfies Definition 4.2. Observe that L-RAMBO and RAMBO have the same external signatures.

We denote by MSG_{RA} the set of messages in the channel automata of RAMBO and by MSG_{LR} the set of messages in the channel automata of L-RAMBO.

We define the simulation relation R to map:

- (α) a state t of L-RAMBO to a state s of RAMBO so that every “common” state variable has the same value. For example, for node $i \in I$, $t.world_i = s.world_i$, $t.pnum1_i = s.pnum1_i$, $t.cmap_i = s.cmap_i$, etc.
- (β) a message $m = \langle W, D, v, t, cm, pns, pnr \rangle \in Channel_{i,j}.MSG_{LR}$ to a message $m' = \langle W, v, t, cm, pns, pnr \rangle \in Channel_{i,j}.MSG_{RA}$ so that: (i) $m.v = m'.v$, (ii) $m.t = m'.t$, (iii) $m.cm = m'.cm$, (iv) $m.pns = m'.pns$, (v) $m.pnr = m'.pnr$, and (vi) $m.W = m'.W$.

Recall that the difference between the two algorithms is in the *Reader–Writer* automata. Therefore, in order to show that L-RAMBO simulates RAMBO, we focus only on the transitions related to the *Reader–Writer* automata. (We remind the reader that the detailed code of RAMBO is given in Appendix.) We now show that R satisfies Definition 4.2:

- (1) If t is an initial state of L-RAMBO, then there exists an initial state s of RAMBO such that $s \in R(t)$, since all common state variables have the same initial values. For example, $\forall i \in I$, $t.world_i = \emptyset = s.world_i$, $t.pnum1_i = 0 = s.pnum1_i$, etc.
- (2) Suppose t and s are reachable states of L-RAMBO and RAMBO respectively such that $s \in R(t)$ and that (t, π, t') . We show that there exists a state $s' \in R(t')$ such that there is an execution fragment of RAMBO that has the same trace as π .
 - (a) If $\pi = \text{send}(m)_{i,j}$, $i, j \in I$, where $m = \langle W, D, v, t, cm, pns, pnr \rangle$ then let s' be such that $(s, \text{send}(m')_{i,j}, s')$, where $m' = \langle W, v, t, cm, pns, pnr \rangle$. Both actions have an empty trace, since $\text{send}(m)_{i,j}$ (resp. $\text{send}(m')_{i,j}$) is considered internal with respect to the composition of automata that comprises automaton L-RAMBO (resp. RAMBO).

From the code of L-RAMBO, we have that $m = \langle W, D, v, t, cm, pns, pnr \rangle$ is placed in the $Channel_{i,j}.MSG_{LR}$, where $m.W = t.world_i$, $m.v = t.value_i$, $m.cm = t.cmap_i$, $m.pns = t.pns_i$, and $m.pnr = t.pnum2(j)_i$. From the state correspondence, since $\text{send}(m)_{i,j}$ is enabled, $\text{send}(m')_{i,j}$ is also enabled. Furthermore, $\text{send}(m')_{i,j}$ places $m' = \langle W, v, t, cm, pns, pnr \rangle$ in the $Channel_{i,j}.MSG_{RA}$ where $m'.W = s.world_i$, $m'.v = s.value_i$, $m'.cm = s.cmap_i$, $m'.pns = s.pns_i$, and $m'.pnr = s.pnum2(j)_i$. From the state correspondence of R for t and s , we conclude that the message correspondence of R for t' and s' is preserved. Also, the state correspondence for t' and s' is preserved, since $t'.pnum1_i = t.pnum1_i + 1 = s'.pnum1_i + 1 = s'.pnum1_i$ and all other common variables remain unchanged.

- (b) If $\pi = \text{recv}(m)_{j,i}$, $i, j \in I$, where $m = \langle W, D, v, t, cm, pns, pnr \rangle$. By Lemma 4.1, $m \in s.h\text{-MSG}$. Let s' be such that $(s, \text{recv}(m')_{j,i}, s')$, where $m' = \langle W, v, t, cm, pns, pnr \rangle$. Both actions have empty trace, since $\text{recv}(m)_{j,i}$ (resp. $\text{recv}(m')_{j,i}$) is considered internal in the composition of automata that comprises automaton L-RAMBO (resp. RAMBO).

Now, since $m = \langle W, D, v, t, cm, pns, pnr \rangle$ was in $\text{Channel}_{i,j}.MSG_{LR}$, by the message correspondence of R , message $m' = \langle W, v, t, cm, pns, pnr \rangle$ is in $\text{Channel}_{i,j}.MSG_{RA}$ and m and m' satisfy the mapping defined above. By inspection of the code of the algorithms, the only difference between update performed by L-RAMBO and RAMBO is that L-RAMBO updates the departed_i set. The the remaining common variables undergo identical updates; hence the state correspondence for t' and s' is preserved. Also, the message correspondence is not affected, since no messages are sent.

- (c) If $\pi = \text{send}(m)_{i,j}$, $i, j \in I$, where $m = \text{leave}$, then let s' be such that $s = s'$ and RAMBO does not perform any action. The action π in L-RAMBO has an empty trace, since it is considered internal in the composition of automata that comprises automaton L-RAMBO. By the code of this action, only the leave-world_i variable is updated. The remaining common state variables of L-RAMBO are unchanged. Since RAMBO does not perform a step, its state variables are not modified. Hence, the state correspondence for t' and s' is preserved.
- (d) If $\pi = \text{recv}(m)_{j,i}$, $i, j \in I$, where $m = \text{leave}$. By Lemma 4.1, $m \in s.h\text{-MSG}$. Let s' be such that $s = s'$ and RAMBO does not perform any action. The action π in L-RAMBO has an empty trace, since it is considered internal in the composition of automata that comprises automaton L-RAMBO. By the code of this action, only the departed_i set is updated. The remaining common state variables of L-RAMBO are unchanged. Since RAMBO does not perform a step, its state variables are not modified. Hence, the state correspondence for t' and s' is preserved.
- (e) If $\pi = \text{leave}_i$, $i \in I$, then let s' be such that (s, leave_i, s') . By examination of the code, the state variable failed_i is set to *true* by both L-RAMBO and RAMBO. The remaining common state variables are not changed. (Algorithm L-RAMBO initializes leave-world_i with the identifiers found in the set $\{\text{world}_i - \text{departed}_i - \{i\}\}$.) Since, $s'.\text{failed}_i = \text{true} = t'.\text{failed}_i$, and all remaining common state variables are unchanged, the state correspondence for t' and s' is preserved.
- (f) If π is not one of the above actions, then we choose s' such that (s, π, s') . That is, we simulate the same action. By inspection of the code of the algorithms, it follows that any state change after π has occurred is identical for both algorithms, hence the state correspondence and message correspondence given by R is preserved.

Therefore, R is a simulation mapping from L-RAMBO to RAMBO per Definition 4.2. Thus, L-RAMBO simulates RAMBO. Since RAMBO implements atomic objects [1,2], so does L-RAMBO. \square

Finally, observe from the preconditions of $\text{send}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$ that if $i \in \text{departed}_j$, then send is disabled, i.e., once a node j learns that another node i left the system, j stops gossiping to i . The correctness of the leave service is as follows: if node i is placed in the departed set of node j , then i has indeed departed the service.

Theorem 4.5. *For all states s of an execution of algorithm L-RAMBO, for any $i, j \in I$, if $i \in s.\text{departed}_j$ then $i \in s.\text{departed}_i$.*

Proof. The proof is by induction on the length of the execution. The base case holds trivially, since all sets are empty in the initial state. Assume that the claim of the lemma holds up to state s , and consider step (s, π, s') . Consider the case where $i \in s.\text{departed}_j$, $i, j \in I$. By inspection of the code (see Fig. 3), we see that no identifier is ever removed from departed . Hence, we have that $i \in s'.\text{departed}_j$, and by inductive hypothesis, $i \in s.\text{departed}_i$, and hence $i \in s'.\text{departed}_i$ as desired. Let us now consider the more interesting case where $i \notin s.\text{departed}_j$ and $i \in s'.\text{departed}_j$. That is, the case where j adds i in departed_j during step (s, π, s') . By examination of the code, there are only two actions that would make the above case possible:

- $\pi = \text{recv}(m)_{i,j}$, $i, j \in I$, where $m = \text{leave}$. By Lemma 4.1, $m \in s.h\text{-MSG}$. Hence, there is a preceding $\text{send}(\text{leave})_{i,j}$ event. By inspection of the code, we observe that by the time the preconditions of the output $\text{send}(\text{leave})_{i,j}$ are satisfied, the following is already true: $i \in \text{departed}_i$. Hence, we have that $i \in s'.\text{departed}_i$ as desired.
- $\pi = \text{recv}(m)_{k,j}$, $k, j \in I$ ($k \neq j$), where $m = \langle W, D, v, t, cm, pnr, pns \rangle$, and $i \in D$. By Lemma 4.1, $m \in s.h\text{-MSG}$. Hence, there is a preceding $\text{send}(\langle W, D, v, t, cm, pnr, pns \rangle)_{k,j}$ event, where $i \in D$. By the

preconditions of this action, we have that $i \in \hat{s}.departed_k$ for some state $\hat{s} < s$. By our inductive hypothesis, we have that $i \in \hat{s}.departed_i$ and hence $i \in s'.departed_i$ as desired.

This completes the proof. \square

5. RAMBO with graceful leave and incremental gossip

Now we present, and prove correct, our final algorithm, called LL-RAMBO (Long-Lived RAMBO). The algorithm is obtained by incorporating incremental gossip into L-RAMBO, so that the size of gossip messages is controlled by eliminating redundant information. In L-RAMBO (resp. RAMBO) the gossip messages contain sets corresponding to the sender's *world* and *departed* (resp. *world*) state variables at the time of sending (Fig. 3). As new nodes join the system and as participants leave the system, the cardinality of these sets grows without bound, rendering RAMBO and L-RAMBO impractical for implementing long-lived objects. The LL-RAMBO algorithm addresses this issue. The challenge here is to ensure that only the certifiably redundant information is eliminated from the messages, while tolerating message loss and reordering.

Specification of LL-RAMBO. We specify the algorithm by modifying the code of L-RAMBO. In Fig. 3, these modifications are annotated with a parenthesized asterisk ((*)). The lines annotated with [h] in Fig. 3 deal with history variables (that are used only in the proof of correctness). The new gossip protocol allows node i to gossip the information in the sets $world_i$ and $departed_i$ incrementally to each node $j \in world_i - departed_i$. Following j 's acknowledgment,¹ node i never again includes this information in gossip messages sent to j , but will include new information that i has learned since the last acknowledgment by j .

To describe the incremental gossip idea in more detail, we consider an exchange of a gossip messages between nodes i and j , where i is the sender and j is the receiver. The sets *world* and *departed* are managed independently and similarly, and we illustrate incremental gossip using just the set *world*. First, we define new data types. Let an *incremental gossip identifier* be the tuple $\langle w-known, d-known, w-unack, d-unack, p-ack \rangle$, where *w-known*, *d-known*, *w-unack*, and *d-unack* are finite subsets of I , and *p-ack* is a natural number. Let *IG* denote the set of all *incremental gossip identifiers*. Finally, let *IGMap* be the set of *incremental gossip maps*, defined as the set of mappings $I \rightarrow IG$. We extend the state of the *Reader-Writer* _{i} automaton with $ig_i \in IGMap$. Node i uses a $ig(j)_i$ tuple to keep track of the knowledge it has about the information already in possession of, and currently being propagated to, node j (see Fig. 3). Specifically, for each $j \in world_i$, $ig(j)_i.w-known$ is the set of node identifiers that i is assured is a subset of $world_j$, $ig(j)_i.w-unack$ is the set of node identifiers, a subset of $world_i$, that j needs to acknowledge. The components $ig(j)_i.d-known$ and $ig(j)_i.d-unack$ are defined similarly for the *departed* set. Lastly, $ig(j)_i.p-ack$ is the phase number of i when the last acknowledgment from j was received. Initially each of these sets is empty, and *p-ack* is zero for each $ig(j)_i$ with $j \in I$.

Node j *acknowledges* a set of identifiers by including this set in the gossip message, or by sending a phase number of i such that node i can deduce that node j has received this set of identifiers in some previous message from i to j . Messages that include i 's phase number that is larger than $ig(j)_i.p-ack$ are referred to as *fresh* or *acknowledgment* messages; otherwise they are referred to as *stale* messages. (We will return to this point in the discussion that follows.)

In RAMBO, once node i learns about node j , it can gossip to j at any time. We now examine the $send((W, D, v, t, cm, pns, pnr))_{i,j}$ action. The world component, W , is set to the difference of $world_i$ and the information that i knows that j has, $ig(j)_i.w-known$, at the time of the send. The remaining components of the gossip message are the same as in L-RAMBO. The effect of the *send* action causes the phase number of the sender to increase; this ensures that each message sent is labeled with a unique phase number of the sender.

Now we examine the $recv((W, D, v, t, cm, pns, pnr))_{i,j}$ action at j (note that we switch i and j relative to the code in Fig. 3 to continue referring to the interaction of the sender i and receiver j). The component W contains a subset of node identifiers from j 's *world*. Hence, W is always used to update $world_j$, $ig(i)_j.w-known$, and $ig(i)_j.w-unack$. The update of $world_j$ is identical to that in L-RAMBO. By definition, $ig(i)_j.w-known$ is the set of node identifiers that j is assured that i has; hence we update it with information in W . Similarly, by definition $ig(i)_j.w-unack$ is the set of node identifiers that j is waiting for i to acknowledge. It is possible that i has learned some

¹ We note that this is not an explicit acknowledgment of a message, but some future message that contains information about what that node learned.

Table 4
Additional state variables in the specification of LL-RAMBO and history variables used in the proof of correctness

Additional state variables used in the specification of LL-RAMBO.	
<i>IG</i>	A set of incremental gossip records. For each record $ig \in IG$, we define the following fields: (i) <i>w-known</i> , (ii) <i>w-unack</i> , (iii) <i>d-known</i> , (iv) <i>d-unack</i> , where each is $\subseteq I$ and each is initially set to \emptyset , and (v) <i>p-ack</i> $\in \mathbb{N}$. These records are used to keep track of knowledge exchange between two nodes in I .
<i>IGMap</i>	A set of incremental gossip maps, defined as a set of mappings from \mathbb{N} to IG .
<i>hs-pack</i>	A mapping from $I \times I \times \mathbb{N}$ to \mathbb{N} . A history variable that records value of $ig(j)_i.p\text{-ack}$ during each send event from i to j .

History variables used in the proof of LL-RAMBO correctness. Each of the following is a mapping from $I \times I \times \mathbb{N}$ to $2^I \cup \{\perp\}$.	
<i>hsent-W</i> and <i>hsent-D</i>	<i>hsent-W</i> records the <i>world</i> component of each message sent by i . <i>hsent-D</i> is defined analogously for the <i>departed</i> set.
<i>hrecv-W</i> and <i>hrecv-D</i>	<i>hrecv-W</i> records the <i>world</i> component of a each message received by i . <i>hrecv-D</i> is defined analogously for the <i>departed</i> set.
<i>hs-world</i> and <i>hs-departed</i>	<i>hs-world</i> records value of the <i>world</i> _{i} during each send event. <i>hs-departed</i> is defined analogously for the <i>departed</i> set.
<i>hs-w-known</i> and <i>hs-d-known</i>	<i>hs-w-known</i> records value of the $ig(j)_i.w\text{-known}$ at each time a message is sent from i to j . <i>hs-d-known</i> is defined analogously for the <i>departed</i> set.
<i>hs-w-unack</i> and <i>hs-d-unack</i>	<i>hs-w-unack</i> records value of the $ig(j)_i.w\text{-unack}$ at each time a message is sent from i to j . <i>hs-d-unack</i> is defined analogously for the <i>departed</i> set.

or all of this information from other nodes, and it is now a part of W ; hence we remove any identifiers in W that are also in $ig(i)_j.w\text{-unack}$ from $ig(i)_j.w\text{-unack}$; these identifiers do not need further acknowledgment.

What happens next in the effect of `rcv` depends on the value of pnr (the phase number that i believes j to be in). First, if $pnr \leq ig(i)_j.p\text{-ack}$, this means that this message is a stale message, since there must have been a prior message from j to i that included phase number of j higher or equal to pnr . Hence, no updates take place. Second, if $pnr > ig(i)_j.p\text{-ack}$, this message is considered to be an acknowledgment message. By definition, $ig(i)_j.p\text{-ack}$ contains the phase number of j when the last acknowledgment from i was received. Following the last acknowledgment, the phase number of j was incremented, $ig(i)_j.p\text{-ack}$ was assigned the new value of the phase number of j , and lastly the new set of identifiers to be propagated was recorded. Since node i replied to j with a phase number larger than $ig(i)_j.p\text{-ack}$, it means that j and i exchanged messages where i learned about the new phase number of j ; by the same token i also learned the information included in these messages. (We show formally that $ig(i)_j.w\text{-unack}$ is always a subset of each message component W that is sent to i by j .) Hence, it is safe for j to assume that i at least received the information in $ig(i)_j.w\text{-unack}$, and to add it to $ig(i)_j.w\text{-known}$.

Since the choice of i and j is arbitrary, gossip from j to i is defined identically.

Brief descriptions of variables introduced in this section are in Table 4. The remaining variables are given in Tables 2 and 3, and the data types are given in Table 1.

Atomicity of LL-RAMBO. We show that any trace of LL-RAMBO is a trace of L-RAMBO, and thus a trace of RAMBO. We start by defining the remaining history variables used in the proofs. These variables are annotated in Fig. 3 with a $[h]$ symbol.

- For every tuple $\langle m, i, j \rangle \in h\text{-msg}_i$, where $m = \langle W, D, v, t, cm, pns, pnr \rangle$ and $pns = p$, the history variable $hsent\text{-}W(i, j, p)$ is a mapping from $I \times I \times \mathbb{N}$ to $2^I \cup \{\perp\}$. This variable records the *world* component of the message, W , when i sends message m to j , and i 's phase number is p . Similarly, we define a derived history variable $hsent\text{-}D(i, j, p)$, a mapping from $I \times I \times \mathbb{N}$ to $2^I \cup \{\perp\}$. This history variable records the *departed* component of the message, D , when i sends message m to j , and i 's phase number is p .

Now we list history variables used to record information for each $\text{send}(\langle W, D, v, t, cm, pns, pnr \rangle)_{i,j}$ event.

- Each of the following variables is a mappings from $I \times I \times \mathbb{N}$ to $2^I \cup \{\perp\}$. $hs\text{-}world(i, j, pns)$ records the value of $world_i$, $hs\text{-}departed(i, j, pns)$ records the value of $departed_i$, $hs\text{-}w\text{-}known(i, j, pns)$ records the value of $ig(j)_i.w\text{-}known$, $hs\text{-}d\text{-}known(i, j, pns)$ records the value of $ig(j)_i.d\text{-}known$, $hs\text{-}w\text{-}unack(i, j, pns)$ records the value of $ig(j)_i.w\text{-}unack$, and $hs\text{-}d\text{-}unack(i, j, pns)$ records the value of $ig(j)_i.d\text{-}unack$.
- $hs\text{-}pack(i, j, pns)$ is a mapping from $I \times I \times \mathbb{N}$ to \mathbb{N} . It records the value of $ig(j)_i.p\text{-}ack$.

The last history variables record information in messages at each $\text{recv}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$ event.

- Each of the following is a mapping from $I \times I \times \mathbb{N}$ to $2^I \cup \{\perp\}$. $hrecv\text{-}W(j, i, pns)$ records the component W (world) and $hrecv\text{-}D(j, i, pns)$ records the component D (departed).

Similarly as in Section 4, we define a history variable $h\text{-}MSG$ that keeps track of messages sent by *Reader–Writer* automata.

Lemma 5.1. *In any execution of LL-RAMBO, if m is a message received by node i in a $\text{recv}(m)_{i,j}$ event, then $\langle m, j, i \rangle \in h\text{-}MSG$, and $m \in \{\langle W, D, v, t, cm, pns, pnr \rangle, \text{leave}, \text{join}\}$, where $\langle W, D, v, t, cm, pns, pnr \rangle \in I \times I \times V \times T \times Cmap \times \mathbb{N} \times \mathbb{N}$.*

Proof. This proof is identical to that of Lemma 4.1, since the format of messages sent by *Reader–Writer* automata in LL-RAMBO is as in L-RAMBO. \square

We continue by showing the properties of messages delivered by *Reader–Writer* processes, which are considered as fresh responses — based on the examination of the message component that is the phase number of the receiver. Specifically, if the message is deemed as an acknowledgment message, then the information that is expected to be acknowledged by the receiver is always a subset of the information included in the message from the sender.

Lemma 5.2. *Consider a step $\langle s, \pi, s' \rangle$ of an execution α of LL-RAMBO, where $\pi = \text{recv}(\langle W, D, v, t, cm, p_j, p_i \rangle)_{j,i}$ for $i, j \in I$, and $p_i > s.ig(j)_i.p\text{-}ack$. Then, (a) $s.ig(j)_i.p\text{-}ack = s.hs\text{-}pack(i, j, p_i)$, (b) $s.ig(j)_i.w\text{-}unack \subseteq s.hs\text{-}w\text{-}unack(i, j, p_i)$, and (c) $s.ig(j)_i.d\text{-}unack \subseteq s.hs\text{-}d\text{-}unack(i, j, p_i)$.*

Proof. We prove the three parts separately:

Part (a). Assume for contradiction that $p_i > s.ig(j)_i.p\text{-}ack$ and $s.hs\text{-}pack(i, j, p_i) \neq s.ig(j)_i.p\text{-}ack$. By the code of the algorithm and the monotonicity of $pnum1$, the only possibility is such that $s.hs\text{-}pack(i, j, p_i) < s.ig(j)_i.p\text{-}ack$. This suggests that there must be a receive event $\text{recv}(\langle W, D, v, t, cm, p_j, p \rangle)_{j,i}$, where $p > s.hs\text{-}pack(i, j, p_i)$ (hence $hrecv\text{-}W(j, i, p)$ and $hrecv\text{-}D(j, i, p)$ are defined) that resulted in the value of $s.ig(j)_i.p\text{-}ack$. By the code of the $\text{recv}_{j,i}$ action, for $i, j \in I$, $ig(j)_i.p\text{-}ack$ is assigned the phase number of i that i has during the receive event. Hence, by the phase number paradigm, we have that $s.ig(j)_i.p\text{-}ack \geq p_i$, which contradicts our initial assumption.

Part (b). From part (a) we have that $hs\text{-}pack(i, j, p_i) = s.ig(j)_i.p\text{-}ack$. From the code, it follows that if $ig(j)_i.p\text{-}ack$ does not change, then the membership of the set $ig(j)_i.w\text{-}unack$ can only be reduced (the “if $pnr > ig(j)_i.p\text{-}ack$ then” statement in not executed). Therefore, $s.ig(j)_i.w\text{-}unack \subseteq s.hs\text{-}w\text{-}unack(i, j, p_i)$.

Part (c). Similar to part (b). From part (a), we have that $hs\text{-}pack(i, j, p_i) = s.ig(j)_i.p\text{-}ack$. From the code, it follows that if $ig(j)_i.p\text{-}ack$ does not change, then the membership of the set $ig(j)_i.d\text{-}unack$ can only be reduced (the “if $pnr > ig(j)_i.p\text{-}ack$ then” statement in not executed). Therefore, $s.ig(j)_i.d\text{-}unack \subseteq s.hs\text{-}d\text{-}unack(i, j, p_i)$. \square

We now state and prove four invariants that lead to the proof of atomicity of LL-RAMBO. The first invariant states that a node does not send information to another node that the receiving node does not already possess.

Invariant 1. *For all states s of any execution α of LL-RAMBO:*

$$\langle \langle W, D, v, t, cm, pns, pnr \rangle, i, j \rangle \in s.h\text{-}MSG \Rightarrow W \subseteq s.world_i \wedge D \subseteq s.departed_i.$$

Proof. The proof is by induction on the length of the execution. The base case is trivial, because all sets are empty in the initial state. Assume the invariant holds for state s and consider step (s, π, s') .

- (1) $\pi = \text{recv}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$. By Lemma 5.1, $\langle \langle W, D, v, t, cm, pns, pnr \rangle, j, i \rangle \in h\text{-}MSG$. By the inductive hypothesis and the monotonicity of the sets $world_i$ and $departed_i$, in state s' , the invariant is maintained.

- (2) $\pi = \text{recv}(\text{join})_{j,i}$. By Lemma 5.1, $\langle \text{join}, j, i \rangle \in h\text{-MSG}$. By the effects of this action, $s.\text{world}_i \subseteq s'.\text{world}_i$ and $s.\text{departed}_i = s'.\text{departed}_i$. By the inductive hypothesis and the monotonicity of the sets world_i and departed_i , in state s' , the invariant is maintained.
- (3) $\pi = \text{recv}(\text{leave})_{j,i}$. By Lemma 5.1, $\langle \text{leave}, j, i \rangle \in h\text{-MSG}$. By the effects of this action, $s.\text{world}_i = s'.\text{world}_i$ and $s.\text{departed}_i \subseteq s'.\text{departed}_i$. By the inductive hypothesis and the monotonicity of the sets world_i and departed_i , in state s' , the invariant is maintained.
- (4) $\pi = \text{send}(m)_{i,j}$. By the code of this action $m = \langle W, D, v, t, cm, pns, pnr \rangle$, such that $W = s.\text{world}_i - s.\text{ig}(j)_i.w\text{-known}$ and $D = s.\text{departed}_i - s.\text{ig}(j)_i.d\text{-known}$. Therefore, $W \subseteq s.\text{world}_i$ and $D \subseteq s.\text{departed}_i$. Observe that the effects of this action do not change world_i and departed_i . Hence, $W \subseteq s'.\text{world}_i$ and $D \subseteq s'.\text{departed}_i$. By the inductive hypothesis, the assignment $\langle m, i, j \rangle \in h\text{-MSG}$ maintains the invariant.
- (5) Other actions do not change the variables involved in the invariant. So, by the inductive hypothesis, in state s' , the invariant continues to hold.

This completes the proof. \square

The following invariant states that the information that i expects j to acknowledge does not exceed the information that actually should be acknowledged.

Invariant 2. For all states s of any execution α of LL-RAMBO:

- (a) $\forall i, j \in I : s.\text{ig}(j)_i.w\text{-unack} \subseteq s.\text{world}_i - s.\text{ig}(j)_i.w\text{-known}$.
- (b) $\forall i, j \in I : s.\text{ig}(j)_i.d\text{-unack} \subseteq s.\text{world}_i - s.\text{ig}(j)_i.d\text{-known}$.

Proof. We prove part (a) of the invariant. The proof of part (b) is analogous: the arguments are made on departed related variables instead on world related variables.

The proof is by induction on the length of the execution. The base case is trivial because all sets are empty in the initial state. Assume the invariant (part (a)) holds for state s , and consider step (s, π, s') .

- (1) $\pi = \text{recv}(\langle W, D, v, t, cm, pns, pnr \rangle)_{j,i}$. By Lemma 5.1, $\langle \langle W, D, v, t, cm, pns, pnr \rangle, j, i \rangle \in s.h\text{-MSG}$. We consider 4 cases:
 - (i) $z \in W \wedge z \notin s.\text{world}_i \wedge z \neq j$. This is the first time node i learns about node z , indirectly from j . By the initial value of the $\text{ig}(z)_i$ record and the monotonicity of world_i the invariant (part (a)) is maintained.
 - (ii) $j \notin s.\text{world}_i$. This is the first time i learns about j , directly from j . Similarly to case 1(i), $\text{ig}(j)_i.w\text{-unack}$ and $\text{ig}(j)_i.w\text{-known}$ are initialized to empty. Also, $\text{ig}(j)_i.p\text{-ack}$ is set to zero. Before the first inner “if-statement”, $s'.\text{world}_i = s.\text{world}_i \cup W$ (note that j is now in $s'.\text{world}_i$). Also, $s'.\text{ig}(j)_i.w\text{-known} = \text{ig}(j)_i.w\text{-known} \cup W$ and $s'.\text{ig}(j)_i.w\text{-unack} = \text{ig}(j)_i.w\text{-unack} - W$. Since $j \notin s.\text{world}_i$, by the code of the send action i never sends messages to j , and hence pnr is zero. Therefore, the “if $pnr > \text{ig}(j)_i.p\text{-ack}$ then” statement is not executed. By the inductive hypothesis (part (a)), the invariant (part (a)) is reestablished.
 - (iii) $j \in \text{world}_i \wedge pnr \leq s.\text{ig}(j)_i.p\text{-ack}$. Since $pnr \leq s.\text{ig}(j)_i.p\text{-ack}$, this implies that node i has learned and communicated with node j in some earlier step of the execution. By the effects of this action, $s'.\text{world}_i = s.\text{world}_i \cup W$, $s'.\text{ig}(j)_i.w\text{-known} = s.\text{ig}(j)_i.w\text{-known} \cup W$, and $s'.\text{ig}(j)_i.w\text{-unack} = s.\text{ig}(j)_i.w\text{-unack} - W$. Then,

$$\begin{aligned} s'.\text{world}_i - s'.\text{ig}(j)_i.w\text{-known} &= (s.\text{world}_i \cup W) - (s.\text{ig}(j)_i.w\text{-known} \cup W) \\ &= s.\text{world}_i - s.\text{ig}(j)_i.w\text{-known} - W. \end{aligned}$$

By the inductive hypothesis (part (a)),

$$s.\text{ig}(j)_i.w\text{-unack} - W \subseteq s.\text{world}_i - s.\text{ig}(j)_i.w\text{-known} - W.$$

Therefore,

$$s'.\text{ig}(j)_i.w\text{-unack} \subseteq s'.\text{world}_i - s'.\text{ig}(j)_i.w\text{-known}.$$

Thus, the invariant (part (a)), in state s' , is maintained.

- (iv) $j \in \text{world}_i \wedge pnr > s.\text{ig}(j)_i.p\text{-ack}$. Using similar arguments as in case 1(iii), we have that up to the “if $pnr > \text{ig}(j)_i.p\text{-ack}$ then” statement the invariant (part (a)) holds. Since $pnr > s.\text{ig}(j)_i.p\text{-ack}$, the “if-statement” is executed, and by its effects we have $s'.\text{ig}(j)_i.w\text{-unack} = s'.\text{world}_i - s'.\text{ig}(j)_i.w\text{-known}$. Hence, the invariant (part (a)) is reestablished.
- (2) $\pi = \text{recv}(\text{join})_{j,i}$. By Lemma 5.1, $\langle \text{join}, j, i \rangle \in s.h\text{-MSG}$. By the code of this action $s.\text{world}_i \subseteq s'.\text{world}_i$. If this is the first time node i learns about node j . then by the initial assignment of the $\text{ig}(j)_i$ record and the

monotonicity of $world_i$, in state s' , the invariant (part (a)) holds. On the other hand, if $j \in s.world_i$, then by the inductive hypothesis (part (a)) and the monotonicity of $world_i$, in state s' , the invariant (part (a)) holds.

- (3) Other actions do not change the variables involved in the invariant (part (a)). So, by the inductive hypothesis (part (a)), in state s' , the invariant (part (a)) continues to hold.

This completes the proof of part (a). As mentioned at the beginning of the proof, part (b) is shown in a similar manner. \square

In the next invariant, we use history variables to show that if i sends a message to j when it believes that j 's phase number is p , then the information that i expects j to acknowledge does not exceed the information included in that message.

Invariant 3. For all states s of any execution α of LL-RAMBO:

- (a) $\langle\langle W, D, v, t, cm, p, pnr \rangle, i, j \rangle \in s.h\text{-MSG} \Rightarrow s.hs\text{-}w\text{-unack}(i, j, p) \subseteq W$
 (b) $\langle\langle W, D, v, t, cm, p, pnr \rangle, i, j \rangle \in s.h\text{-MSG} \Rightarrow s.hs\text{-}d\text{-unack}(i, j, p) \subseteq D$.

Proof. The proof is by induction on the length of the execution. The base case is trivial because all sets are empty in the initial state. Assume the invariant holds for state s , and consider step (s, π, s') .

- (1) $\pi = \text{send}(\langle\langle W, D, v, t, cm, pns, pnr \rangle\rangle)_{i,j}$. By the effects of this action we have that $W = s'.world_i - s'.ig(j)_i.w\text{-known}$ and $s'.hs\text{-}w\text{-unack}(i, j, p) = s'.ig(j)_i.w\text{-unack}$ is defined. Also, $D = s'.departed_i - s'.ig(j)_i.d\text{-known}$ and $s'.hs\text{-}d\text{-unack}(i, j, p) = s'.ig(j)_i.d\text{-unack}$ is defined. By **Invariant 2** we have that $s'.ig(j)_i.w\text{-unack} \subseteq s'.world_i - s'.ig(j)_i.w\text{-known}$, and $s'.ig(j)_i.d\text{-unack} \subseteq s'.departed_i - s'.ig(j)_i.d\text{-known}$. Thus, by substitution we get that $s'.hs\text{-}w\text{-unack}(i, j, p) \subseteq W$ and $s'.hs\text{-}d\text{-unack}(i, j, p) \subseteq D$, as desired. Therefore, by the inductive hypothesis, in state s' , the invariant is reestablished.
- (2) Other actions do not change the variables involved in the invariant. So, by the inductive hypothesis, in state s' , the invariant continues to hold.

This completes the proof. \square

We use **Lemmas 5.1** and **5.2**, and **Invariants 1–3** to show the key **Invariant 4** for the atomicity of LL-RAMBO. Here, we show that node i never overestimates the knowledge possessed by j .

Invariant 4. For all states s of any execution α of LL-RAMBO:

- (a) $\forall i, j \in I : s.ig(j)_i.w\text{-known} \subseteq s.world_j$,
 (b) $\forall i, j \in I : s.ig(j)_i.d\text{-known} \subseteq s.departed_j$.

Proof. We prove part (a) of the invariant. The proof of part (b) is analogous: the arguments are made on *departed* related variables instead on *world* related variables.

The proof is by induction on the length of the execution. The base case is trivial because all sets are empty in the initial state. Assume the invariant (part (a)) holds for state s and consider step (s, π, s') (the following discussion involves only part (a) of the invariant).

- (1) $\pi = \text{recv}(\langle\langle W, D, v, t, cm, pns, pnr \rangle\rangle)_{*,j}$. By **Lemma 5.1**, $\langle\langle W, D, v, t, cm, pns, pnr \rangle, *, j \rangle \in s.h\text{-MSG}$. The set $s.ig(j)_i.w\text{-known}$, for $i \in I, i \neq j$, is not updated by the effects of this action. Observe that $s.world_j \subseteq s'.world_j$ (for $i = j$ the invariant is trivially maintained). By the inductive hypothesis and the monotonicity of $world_j$, in state s' , the invariant is maintained.
- (2) $\pi = \text{recv}(\text{join})_{*,j}$. By **Lemma 5.1** $\langle\text{join}, *, j \rangle \in s.h\text{-MSG}$. The discussion in this case is identical to that of case 1.
- (3) $\pi = \text{recv}(\langle\langle W, D, v, t, cm, pns, pnr \rangle\rangle)_{j,i}$. By **Lemma 5.1**, $\langle\langle W, D, v, t, cm, pns, pnr \rangle, j, i \rangle \in s.h\text{-MSG}$. We consider 4 cases:
- (i) $z \in W \wedge z \notin s.world_i \wedge z \neq j$. This is the first time node i learns about node z , indirectly from j . By the initial value of the $ig(z)_i$ record and the monotonicity of $world_i$, the invariant is maintained. (This is for all nodes other than j . We consider j in the remaining three subcases.)
- (ii) $j \notin s.world_i$. This is the first time i learns about j , directly from j . Similarly to case 3(i), $ig(j)_i.w\text{-unack}$ and $ig(j)_i.w\text{-known}$ are initialized to empty. Also $ig(j)_i.p\text{-ack}$ is set to zero. Before the “if-statement”, $s'.world_i = s.world_i \cup W$ (note that j is now in $s'.world_i$). Also, $s'.ig(j)_i.w\text{-known} = ig(j)_i.w\text{-known} \cup W$

$(ig(j)_i.w-known = W)$ and $s'.ig(j)_i.w-unack = ig(j)_i.w-unack - W$ ($ig(j)_i.w-unack$ remains empty). Since $j \notin s.world_i$, by the code of the send action i never sent messages to j , and hence pnr is zero. Therefore, the “if $pnr > ig(j)_i.p-ack$ then” statement is not executed. By substitution $s'.ig(j)_i.w-known = W$. From **Invariant 1** we have that $W \subseteq s.world_j = s'.world_j$. Hence $s'.ig(j)_i.w-known \subseteq s'.world_j$, as desired. By the inductive hypothesis, in s' , the invariant is maintained.

- (iii) $j \in s.world_i \wedge pnr \leq s.ig(j)_i.p-ack$. Since $pnr \leq s.ig(j)_i.p-ack$, this implies that node i has learned and communicated with node j in some earlier step of the execution. By the effects of this action $s'.ig(j)_i.w-known = s.ig(j)_i.w-known \cup W$. From **Invariant 1**, we have that $W \subseteq s.world_j = s'.world_j$. Also by the inductive hypothesis $s.ig(j)_i.w-known \subseteq s'.world_j$. Therefore, $s'.ig(j)_i.w-known \subseteq s'.world_j$, as desired. Thus, in s' , the invariant is maintained.
- (iv) $j \in s.world_i \wedge pnr > s.ig(j)_i.p-ack$. Using similar arguments as in case 3(iii) we have that up to the “if $pnr > ig(j)_i.p-ack$ then” $ig(j)_i.w-known = s.ig(j)_i.w-known \cup W$ and $ig(j)_i.w-unack = s.ig(j)_i.w-unack - W$. Since $pnr > s.ig(j)_i.p-ack$ the “if-statement” is executed; by its effects we get $s'.ig(j)_i.w-known = ig(j)_i.w-known \cup ig(j)_i.w-unack$. Recall that $ig(j)_i.w-known = s.ig(j)_i.w-known \cup W$. By **Invariant 1** and the inductive hypothesis, $ig(j)_i.w-known \subseteq s'.world_j$.

Since pnr is contained in the message from j by i , we have that $s.hrecv-W(i, j, pnr)$ must be defined; moreover, by the code of the algorithm, the monotonicity of the $world$ variable, and the definition of $hrecv-W(i, j, pnr)$, we have that $s.hrecv-W(i, j, pnr) \subseteq s'.world_j$. By the properties of the *Channel* automata, we have that $s.hs-w-unack(i, j, pnr)$ must also be defined and its value equals the value of $s.hrecv-W(i, j, pnr)$. By **Invariant 3** we have that $s.hs-w-unack(i, j, pnr) \subseteq s.hs-sent-W(i, j, pnr)$. By **Lemma 5.2(b)** we have that if $pnr > s.ig(j)_i.p-ack$ then $s.ig(j)_i.w-unack \subseteq s.hs-w-unack(i, j, pnr)$. By substitution we have $s.ig(j)_i.w-unack \subseteq s'.world_j$.

Therefore, $s'.ig(j)_i.w-known \subseteq s'.world_j$, as desired. Thus, by the inductive hypothesis, in s' , the invariant is reestablished.

- (4) Other actions do not change the variables involved in the invariant. So, the by inductive hypothesis, in state s' , the invariant continues to hold.

This completes the proof of part (a). As mentioned at the beginning of the proof, part (b) is shown in a similar manner. \square

Finally, we show the atomicity of objects implemented by LL-RAMBO by proving that it simulates L-RAMBO, i.e., by showing that every trace of LL-RAMBO is a trace of L-RAMBO (hence of RAMBO).

Theorem 5.3. *LL-RAMBO implements atomic read/write objects.*

Proof. We show that the LL-RAMBO simulates the L-RAMBO. Specifically, we show that there exists a simulation relation R from LL-RAMBO automaton to L-RAMBO automaton that satisfies **Definition 4.2**. Observe that LL-RAMBO and L-RAMBO have the same external signatures.

We denote by MSG_{LR} the set of messages in the channel automata of L-RAMBO, and by MSG_{LLR} the set of messages in the channel automata of LL-RAMBO.

We define the simulation relation R to map:

(a) a state t of LL-RAMBO to a state s of L-RAMBO so that every “common” state variable has the same value. For example, for node $i \in I$, $t.world_i = s.world_i$, $t.departed_i = s.departed_i$, $t.pnum1_i = s.pnum1_i$, $t.cmap_i = s.cmap_i$, etc.

(b) a message $m = \langle W, D, v, t, cm, pns, pnr \rangle \in Channel_{i,j}.MSG_{LLR}$ to a message $m' = \langle W, D, v, t, cm, pns, pnr \rangle \in Channel_{i,j}.MSG_{LR}$ so that: (i) $m.v = m'.v$, (ii) $m.t = m'.t$, (iii) $m.cm = m'.cm$, (iv) $m.pns = m'.pns$, and $m'.W = hs-world(i, j, pns)$, and lastly (vii) $m.D = hs-departed(i, j, pns) - hs-d-known(i, j, pns)$ and $m'.D = hs-departed(i, j, pns)$. (We assume that the history variables $hs-world(i, j, pns)$ and $hs-departed(i, j, pns)$ are used in L-RAMBO in the similar manner that they are used in LL-RAMBO.)

Recall that the difference between the two algorithms is in the *Reader–Writer* automata. Therefore, in order to show that LL-RAMBO simulates L-RAMBO, we focus only on transitions related to the *Reader–Writer* automata. We now show that R satisfies **Definition 4.2**:

- (1) If t is an initial state of LL-RAMBO, then there exists an initial state s of L-RAMBO such that $s \in R(t)$, since all common state variables have the same initial values. For example, $\forall i \in I, t.world_i = s.world_i = \emptyset, t.departed_i = s.departed_i = \emptyset, t.pnum1_i = s.pnum1_i = 0$, etc. Also, the channels do not contain any messages.
- (2) Suppose t and s are reachable states of LL-RAMBO and L-RAMBO respectively such that $s \in R(t)$ and that (t, π, t') . We show that there exists a state $s' \in R(t')$ such that there is an execution fragment of L-RAMBO that has the same trace as π .

- (a) If $\pi = \text{recv}(m)_{j,i}, i, j \in I$, where $m = \langle W, D, v, t, cm, p, pnr \rangle$. By Lemma 5.1, $m \in s.h\text{-MSG}$. Let s' be such that $(s, \text{recv}(m')_{j,i}, s')$. Both actions have empty trace, since $\text{recv}(m)_{j,i}$ (resp. $\text{recv}(m')_{j,i}$) is considered internal with respect to the composition of automata that comprises automaton LL-RAMBO (resp. L-RAMBO).

Now, since $m = \langle W, D, v, t, cm, p, pnr \rangle$ was in $Channel_{j,i}.MSG_{LLR}$, by the message correspondence of R , message $m' = \langle W, D, v, t, cm, p, pnr \rangle$ is in $Channel_{j,i}.MSG_{LR}$, and m and m' have the mapping defined above. By inspection of the code of both algorithms, it follows that besides state variables $world_i$ and $departed_i$, all other common state variables are updated identically, as the information obtained by messages m and m' is the same for these variables. Also, the message correspondence is not affected, since no messages are sent.

Hence we focus on showing that $t'.world_i = s'.world_i$ and $t'.departed_i = s'.departed_i$. From the code of LL-RAMBO, we have that $t'.world_i = t.world_i \cup m.W$ and $t'.departed_i = t.departed_i \cup m.D$. Since m was received (and hence removed from the channel), by the properties of the channel, there was a send event that placed m in the channel so that $m.W = t.hs\text{-sent-}W(j, i, p) = t.hs\text{-world}(j, i, p) - t.hs\text{-w-known}(j, i, p)$ and $m.D = t.hs\text{-sent-}D(j, i, p) = t.hs\text{-departed}(j, i, p) - t.hs\text{-d-known}(j, i, p)$. From Invariant 4 and the monotonicity of variables $world$ and $departed$, we have that $hs\text{-w-known}(j, i, p) \subseteq t.world_i$ and $hs\text{-d-known}(j, i, p) \subseteq t.departed_i$. From this, and the fact that $t.hs\text{-w-known}(j, i, p) \subseteq t.hs\text{-world}(j, i, p)$ and $t.hs\text{-d-known}(j, i, p) \subseteq t.hs\text{-departed}(j, i, p)$, we conclude that $t'.world_i = t.world_i \cup t.hs\text{-world}(j, i, p)$ and $t'.departed_i = t.departed_i \cup t.hs\text{-departed}(j, i, p)$. But by the state and message correspondence of R , we have that $t.world_i = s.world_i, t.departed_i = s.departed_i, m'.W = t.hs\text{-world}(j, i, p)$, and $m'.D = t.hs\text{-departed}(j, i, p)$. Hence, $t'.world_i = s.world_i \cup m'.W = s'.world_i$ and $t'.departed_i = s.departed_i \cup m'.D = s'.departed_i$, as desired.

- (b) If $\pi = \text{send}(m)_{i,j}, i, j \in I$, where $m = \langle W, D, v, t, cm, p, pnr \rangle$. Let s' be such that $(s, \text{send}(m')_{i,j}, s')$. Both actions have empty trace, since $\text{send}(m)_{i,j}$ (resp. $\text{send}(m')_{i,j}$) is considered internal in the composition of automata that comprises automaton LL-RAMBO (resp. L-RAMBO).

From the code of LL-RAMBO, we have that $m = \langle W, D, v, t, cm, pns, pnr \rangle$ is placed in $Channel_{i,j}.MSG_{LLR}$, where $m.W = hs\text{-world}(i, j, pns) - hs\text{-w-known}(i, j, pns)$, $m.D = hs\text{-departed}(i, j, pns) - hs\text{-d-known}(i, j, pns)$, $m.v = t.value_i, m.cm = t.cmap_i, m.pns = t.pnum1_i$, and $m.pnr = t.pnum2(j)_i$. From the state correspondence, since $\text{send}(m)_{i,j}$ is enabled, $\text{send}(m')_{i,j}$ is also enabled. Furthermore, $\text{send}(m')_{i,j}$ places $m' = \langle W, D, v, t, cm, pns, pnr \rangle$ in $Channel_{i,j}.MSG_{LR}$ where $m'.W = hs\text{-world}(i, j, pns), m'.D = hs\text{-departed}(i, j, pns), m'.v = s.value_i, m'.cm = s.cmap_i, m'.pns = s.pnum1_i$, and $m'.pnr = s.pnum2(j)_i$. From the state correspondence of R for t and s , we conclude that the message correspondence of R for t' and s' is preserved. Also, the state correspondence for t' and s' is preserved, since $t'.pnum1_i = t.pnum1_i + 1 = s.pnum1_i + 1 = s'.pnum1_i$ and all other common variables remain unchanged.

- (c) If π is an action besides $\text{send}(m)_{i,j}$ or $\text{recv}(m)_{j,i}$, then we choose s' such that (s, π, s') . That is, we simulate the same action. By inspection of the code of the algorithms, it follows that any state change after π occurred is identical for both algorithms; hence the state correspondence and message correspondence given by R is preserved.

Therefore, R is a simulation mapping from LL-RAMBO to L-RAMBO per Definition 4.2. By Theorem 4.3, any trace of LL-RAMBO is a trace of L-RAMBO. Since L-RAMBO implements atomic read/write objects per Theorem 4.4, so does LL-RAMBO.

6. LL-RAMBO performance

In this section, we analyze the performance of LL-RAMBO by giving a conditional analysis of read and write operation latency and assessing improvement in communication for specific scenarios.

6.1. Conditional read/write operation latency analysis

A conditional analysis of RAMBO read and write operation latency is presented in [2]. Here, we show that under the same conditions, LL-RAMBO has the same operation latency. We start by giving relevant definitions (based on [2] and [30]). Let d denote the maximum message delivery latency. Let d also be the interval at which the gossip messages are sent. An execution with times associated with all events is called a *timed execution*. A timed execution is said to be *admissible* if the following condition holds: “If timed execution ξ is an infinite sequence, then the times of the actions approach ∞ . If ξ is a finite sequence, then in the final state of ξ , every enabled task must be allowed to complete” (see [30]).

Let α be an admissible timed execution, and let α' be a finite prefix of α . Let $\elltime(\alpha')$ denote the time of the last event in α' . We say α is an α' -normal execution if (i) after α' , the local clocks of all automata progress at exactly the rate of real time, (ii) no message sent in α after α' is lost, and (iii) if a message is sent at time t in α and it is delivered, then it is delivered by the time $\max\{t + d, \elltime(\alpha') + d\}$.

LL-RAMBO allows sending of gossip messages at arbitrary times. For the purpose of latency analysis, we restrict the sending pattern: we assume that each automaton sends messages at the first possible time and at regular intervals of d thereafter, as measured on the local clock. Also, non-send locally controlled events occur just once, within time 0 on the local clock.

As with all quorum-based algorithms, operation liveness depends on all the nodes in some quorums remaining alive or not departing. We say that a configuration is *installed* when every member of the configuration has been notified about the configuration. We say that an execution α is (α', e, τ) -configuration-viable if for every installed configuration, there exists a read-quorum, R , and a write-quorum, W , such that no process in $R \cup W$ fails or departs before the maximum of (i) time τ after the next configuration is installed, and (ii) $\elltime(\alpha') + e + \tau$.

We say that execution α satisfies (α', τ) -recon-spacing if after α' , at least time τ elapses between the event that reports a new configuration c ($\text{report}(c)_i$) and any following event that proposes a new configuration ($\text{recon}(c, *)_i$). In other words, after α' , when the system stabilizes, reconfigurations are not too frequent.

Execution α is said to satisfy (α', e) -join-connectivity if after α' , for any two nodes that both joined the system at time $t - e$, they know about each other by time t .

Execution α satisfies $(\alpha', e + \tau)$ -recon-readiness if after α' , every $\text{recon}(c)$ event proposing a new configuration includes a node i in c only if i joined at least time $e + \tau$ ago. This, in conjunction with (α', e) -join-connectivity, ensures that all the nodes in active configurations are aware of each other.

As in [2], we assume that α is an α' -normal execution, satisfying $(\alpha', e, 23d)$ -configuration-viability, $(\alpha', 8d)$ -recon-spacing, (α', e) -join-connectivity, and $(\alpha', e + d)$ -recon-readiness. (See [2] for a further discussion of these assumptions.) With this we show conditionally that read and write operations take no more than $8d$ time.

Theorem 6.1. *Let α be an α' -normal execution of LL-RAMBO satisfying join-connectivity, recon-readiness, recon-spacing, and configuration-viability. Let $t > \elltime(\alpha') + e + d$. Assume i is a node that received a join-ack_i prior to time $t - e - d$, and neither fails nor departs in α until after time $t + 8d$. Then if a read or write operations starts at node i time t , it completes by time $t + 8d$.*

The proof is essentially identical to the proof of Theorem 5.3 in [2]. The key observation is that under the assumed conditions the incremental gossip does not affect the pattern of messages. The only difference is that these messages do not contain information that had been already propagated.

6.2. Communication efficiency analysis

In this section, we illustrate the communication savings attained by LL-RAMBO as compared to RAMBO. The savings are assessed both in terms of gossip message size and number of gossip messages. We present communication analysis of α' -normal timed executions α , where the prefix α' satisfies specific properties. Following the prefix α' , we assume that α is divided into *rounds* of length d , where d is greater than the maximum message latency. The active nodes send gossip messages at the beginning of each round and subsequently receive gossip messages sent during the current round. We assume that in all executions node failure and departures do not disable the quorum systems used by the algorithms.

We observe that given any timed execution α of LL-RAMBO, it is possible to construct an execution $\hat{\alpha}$ of RAMBO that has the same interaction with its environment as α (as follows from Theorem 5.3), but that includes additional

gossip messages and gossip messages with different content, and that excludes leave notification messages that are specific to LL-RAMBO. We also denote by $\hat{\alpha}'$ the prefix of $\hat{\alpha}$ that corresponds to α' . In the rest of this section, we will be using this notation in comparing the communication efficiency of the two algorithms in terms of gossip messaging.

Recall that gossip messages in RAMBO have the format $\langle W, v, t, cm, pns, pnr \rangle$, and in LL-RAMBO the format is $\langle W, D, v, t, cm, pns, pnr \rangle$. We also assume that each node identifiers is γ bits long, and that the values v, t, cm, pns , and pnr , altogether occupy δ bits (a constant) in gossip messages.

Scenario 1: No joins or departures after α' . Here we make the following additional assumptions about α' for LL-RAMBO and $\hat{\alpha}'$ for RAMBO. Consider the following sequence of events, identical for α' and $\hat{\alpha}'$. The service is initialized by some node, called the *creator*. Next, m new participants join the service by sending to the creator m join requests. After the last join request is received, the creator sends a gossip message to each new participant. Once these gossip messages are received, all nodes have active status. At this point, the cardinality of *world* for each node is $n = m + 1$. Now, l nodes decide to leave the system. In α' of LL-RAMBO, these nodes send leave notification messages to all participating nodes. We assume that these messages are delivered. These notification messages are of constant size and are much smaller than any gossip message. Now the cardinality of the *departed* set at each node is l in LL-RAMBO. In the case of RAMBO, l nodes leave by emulating crash failures. The number of the active nodes, a , is $a = n - l$. This concludes α' (respectively $\hat{\alpha}'$), following which no new nodes join the system or no active nodes leave the system. Following α' (respectively $\hat{\alpha}'$), normal timing holds in α (respectively $\hat{\alpha}$), and the active nodes gossip at regular intervals as described above. We now give the result that compares the two algorithms for r rounds of gossip.

Proposition 6.2. *Let α be α' -normal execution of LL-RAMBO and $\hat{\alpha}$ be $\hat{\alpha}'$ -normal execution of RAMBO as defined by Scenario 1. Then:*

- (a) *there are $r \cdot a \cdot l$ fewer gossip messages in α following α' than in $\hat{\alpha}$ following $\hat{\alpha}'$;*
- (b) *the bit complexity of gossip messages in α following α' is smaller by $((r - 1) \cdot n^2 \cdot a \cdot \gamma) + (a \cdot l^2 \cdot \gamma) + (r \cdot a \cdot l \cdot \delta)$ than the bit complexity in $\hat{\alpha}$ following $\hat{\alpha}'$.*

Proof. Nodes participating in RAMBO cannot distinguish nodes that failed from nodes that departed. Therefore, in each round of gossip, RAMBO sends $a \cdot n$ messages (an active node sends gossip messages to all nodes in its *world*). Each gossip message has the size of $n \cdot \gamma + \delta$ bits ($|W| = |\text{world}| = n$). Hence, the gossip message complexity of RAMBO for r rounds is $r \cdot a \cdot n$, and the gossip message bit complexity is $r \cdot a \cdot n \cdot (n \cdot \gamma + \delta)$.

In LL-RAMBO, each node has learned that l nodes departed. Therefore, a^2 gossip messages are sent in each of the r rounds. In the first round, each gossip message has size $(n + l) \cdot \gamma + \delta$ (since nodes do not know what other nodes know). Since all messages are delivered, at the end of the first round, each node knows that every active node has full knowledge of *world* and *departed* sets. Hence, in the following rounds, the gossip message size is δ bits. Therefore, the gossip message complexity of LL-RAMBO for r rounds is $r \cdot a^2$, and the gossip message bit complexity is $(a^2 \cdot (n + l) \cdot \gamma) + (r \cdot a^2 \cdot \delta)$.

Therefore, LL-RAMBO sends $r \cdot a \cdot l$ fewer gossip messages than RAMBO. The reduction for the gossip message bit complexity is $((r - 1) \cdot n^2 \cdot a \cdot \gamma) + (a \cdot l^2 \cdot \gamma) + (r \cdot a \cdot l \cdot \delta)$. \square

If l is zero, then the behavior of these two algorithms is identical. Otherwise, l can be as small as one and as large as $n - 1$, and since $a = n - l$, the savings in gossip messages for LL-RAMBO are between $\Omega(r \cdot n)$ and $O(r \cdot n^2)$. The reduction in bit complexity for LL-RAMBO is between $\Omega(r \cdot n^2)$ and $O(r \cdot n^3)$ (γ and δ are assumed to be constants).

Scenario 2: Steady turnover rate after α' . In this scenario, following the initial segment of an execution, nodes join and leave at a constant rate per round in both RAMBO and LL-RAMBO. Consider the following sequence of events for both systems. As in Scenario 1, the service is initialized by the creator. Then m new participants join the service by sending join requests to the creator. After the last join request is received, the creator sends gossip to all participants. Once these gossip messages are received, all nodes are active. At this point, the cardinality of *world* at each node is $n = m + 1$. This concludes α' of RAMBO (respectively $\hat{\alpha}'$ of LL-RAMBO). Following α' (respectively $\hat{\alpha}'$), normal timing holds in α (respectively $\hat{\alpha}$), and the active nodes gossip at the beginning of each round. In this scenario, nodes join and leave the service at the same rate: during each round following α' (respectively $\hat{\alpha}'$), z new nodes join and z active nodes leave the service. We now give the result that compares the two algorithms for r rounds of gossip.

Proposition 6.3. *Let α be α' -normal execution of LL-RAMBO and $\hat{\alpha}$ be $\hat{\alpha}'$ -normal execution of RAMBO as defined by Scenario 2. Then:*

- (a) *there are $(n - z) \cdot z \cdot ((r - 1) \cdot (r - 2)/2)$ fewer gossip messages in α following α' than in $\hat{\alpha}$ following $\hat{\alpha}'$;*
 (b) *the bit complexity of gossip messages in α following α' for $r > 2$ is smaller by*

$$(n - z) \cdot \left(\gamma \cdot z^2 \cdot ((r - 1) \cdot (r - 2) \cdot (2 \cdot r - 1)/6) \right. \\ \left. + (2 \cdot \gamma \cdot n + \delta) \cdot z \cdot ((r - 1) \cdot (r - 2)/2) + n \cdot \gamma \cdot (n - z) \cdot (r - 2) + n^2 \cdot \gamma \right)$$

than the bit complexity in $\hat{\alpha}$ following $\hat{\alpha}'$. For $r = 2$ the reduction is $(n - z) \cdot n^2 \cdot \gamma$ and for $r = 1$ the bit complexity is the same.

Proof. For the analysis, we assume that the join requests are received by all active nodes by the end of the round in which they were sent, and are received after all gossip messages are sent. On the other hand, we assume that the leave notifications are received (in LL-RAMBO) by active nodes at the end of the following round. These assumptions yield worst case behavior for LL-RAMBO in Scenario 2, since delaying notification messages only affect the gossip communication complexity of LL-RAMBO (in RAMBO no notification messages are sent).

We now analyze the behavior of RAMBO for a round $1 \leq k \leq r$. By the code of RAMBO, we observe that the *world* set grows as nodes join the system. Since gossip messages include the entire *world* set, it follows that the size of the gossip messages is proportional to the size of *world*. In RAMBO, nodes leave by emulating failures. Since nodes cannot detect such failures, each active node attempts to gossip to all members in its *world*, regardless of their failure status. Hence, as new participants join, the number and the size of gossip messages sent in RAMBO increases.

We separately consider the case when $k = 1$. At the beginning of the first round, the *world* of each active node is n (join requests have not been yet received). From this, and the fact that only active nodes gossip (new joined nodes become active after they receive a gossip message), in first round n^2 messages are sent each of size $n \cdot \gamma + \delta$. We now consider $2 \leq k \leq r$. At the beginning of the k 'th round there are $n - z$ active nodes. Observe that up to k 'th round $(k - 1) \cdot z$ active nodes have left the system; also the same number have attempted to join the system, of which only $(k - 2) \cdot z$ became active by k 'th round (it takes at most $2d$ time for a node to become active). Hence, there may be at most up to $n - z$ active nodes and n participating nodes in the system at the beginning of each round k . The *world* set of each active node has cardinality $n + (k - 1) \cdot z$ (since z nodes enter the system in every round). Therefore, in each round $k > 1$, $(n - z) \cdot (n + (k - 1) \cdot z)$ gossip messages are sent, each of size $(n + (k - 1) \cdot z) \cdot \gamma + \delta$.

We now analyze, in a similar fashion, the behavior of LL-RAMBO for a round $1 \leq k \leq r$. By the code of LL-RAMBO, we observe that both the *world* and the *departed* sets grow as the new participants join and active nodes leave the system. However, two nodes that have been active in two consecutive rounds, need only to exchange gossip messages that contain only the identifiers of the z nodes that have departed and the z nodes that have joined the system in the previous round. On the other hand, an active node has to send all current knowledge about *world* and *departed* to the newly joined nodes. Also, note that once a node learns that another node has departed from the system, it does not send any further gossip messages to that node. We separately consider cases where $k = 1$ and $k = 2$. The argument for the first round is identical to the argument of the first round of RAMBO. At the beginning of the second round, there are $n - z$ active nodes (this is because in the first round z nodes departed and the z nodes joined the system, where the newly joined nodes are not active yet). Also the cardinality of *world* variable of each active node is $n + z$ and the cardinality of *departed* variable is zero (since the leave notifications have not been received yet - they will be received at the end of this round). Therefore, in the second round, $(n - z) \cdot (n + z)$ gossip messages are sent, and the total bit complexity of these messages is $(n - z) \cdot \left(n \cdot (z \cdot \gamma + \delta) + z \cdot ((n + z) \cdot \gamma + \delta) \right)$. We now consider $2 \leq k \leq r$. At the beginning of the round k there are $n - z$ active nodes (using the same reasoning as in case of RAMBO for rounds $k \geq 2$). The *world* set of each active node has cardinality $n + (k - 1) \cdot z$ (since z nodes enter the system in every round), and the *departed* set has cardinality $(k - 2) \cdot z$ (since it takes an extra round for the notifications to be received). It follows that in round $k > 2$, $(n - z) \cdot (n + z)$ gossip messages are sent, and the total bit complexity of these messages is $(n - z) \cdot \left(n \cdot (2 \cdot z \cdot \gamma + \delta) + z \cdot ((n + (k - 1) \cdot z + (k - 2) \cdot z) \cdot \gamma + \delta) \right)$. By performing elementary mathematical computations, it follows that for r rounds, LL-RAMBO sends $(n - z) \cdot z \cdot ((r - 1) \cdot (r - 2)/2)$ less gossip messages than RAMBO. Also, it is not difficult to see that in the first round LL-RAMBO gossip bit complexity

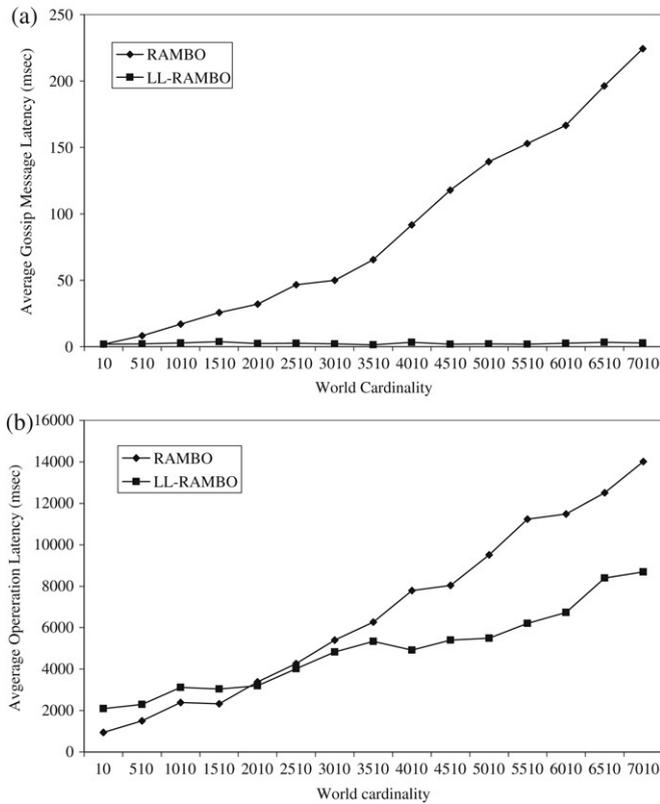


Fig. 4. Preliminary empirical results: (a) gossip message latency, (b) read and write latency.

is the same as RAMBO, and the gossip bit complexity reduction for the second round is $(n - z) \cdot n^2 \cdot \gamma$. Finally, for $r > 2$ we have that the savings in the gossip bit complexity are $(n - z) \cdot (\gamma \cdot z^2 \cdot ((r - 1) \cdot (r - 2) \cdot (2 \cdot r - 1)/6) + (2 \cdot \gamma \cdot n + \delta) \cdot z \cdot ((r - 1) \cdot (r - 2)/2) + n \cdot \gamma \cdot (n - z) \cdot (r - 2) + n^2 \cdot \gamma$. \square

For the following bounds we consider $1 \leq z \leq n - 1$. The savings in gossip messages for LL-RAMBO when $r > 2$ are between $\Omega(r^2 \cdot n)$ and $O(r^2 \cdot n^2)$. When $1 \leq r \leq 2$ there are no savings. The reduction in gossip bit complexity for LL-RAMBO when $r > 1$ is between $\Omega(n^2)$ and $O(r^3 \cdot n^3)$ (γ and δ are assumed to be constants). When $r = 1$ there is no reduction.

7. LL-RAMBO implementation

We developed proof-of-concept implementations of RAMBO and LL-RAMBO on a network-of-workstations. In this section, we present the preliminary experimental results.

Experimental results. We developed the system by manually translating the Input/Output Automata specifications to Java code. To mitigate the introduction of errors during translation, the implementers followed a set of precise rules [3] that guided the derivation of Java code. The platform consists of a Beowulf cluster with ten machines running Linux. The machines are various Pentium processors up to 900 MHz interconnected via a 100 Mbps Ethernet switch. The implementation of the two algorithms share most of the code and all low-level routines, so that any difference in performance is traceable to the distinct *world* and *departed* set management and the gossiping discipline encapsulated in each algorithm.

We are interested in long-lived applications, and we assume that the number of participants grows arbitrarily. Given the limited number of physical nodes, we use majority quorums of these nodes, and we simulate a large number of other nodes that join the system by including such node identifiers in the *world* sets. Using non-existent nodes approximates the behavior of a long-lived system with a large set of participants. However, when using all-to-all gossip

that grows quadratically in the number of participants, it is expected that the differences in RAMBO and LL-RAMBO performance will become more substantial when using a larger number of physical nodes.

The experiment is designed as follows. There are ten nodes that do not leave the system. These nodes perform concurrent read and write operations using a single configuration (that does not change over time), consisting of majorities, i.e., six nodes. Fig. 4 compares: (a) the average latency of gossip messages, and (b) the average latency of read and write operations in RAMBO and LL-RAMBO, as the cardinality of *world* sets grows from 10 to 7010.

LL-RAMBO exhibits substantially better gossip message latency than RAMBO (Fig. 4(a)). In fact, the average gossip latency in LL-RAMBO does not vary noticeably. On the other hand, the gossip latency in RAMBO grows substantially as the cardinality of the *world* sets increases. This is expected, due to the smaller incremental gossip messages of LL-RAMBO, while in RAMBO, the size of the gossip messages is always proportional to the cardinality of the *world* set. LL-RAMBO trades local resources (computation and memory) for smaller and fewer gossip messages. We observe that the read/write operation latency is slightly lower for RAMBO when the cardinality of the *world* sets is small (Fig. 4(b)). As the size of the *world* sets grows, the operation latency in LL-RAMBO becomes substantially better than in RAMBO.

We close this section with the following remarks. First, the experiment is designed using a single configuration that does not change over time. Note that reconfiguration creates additional network traffic; however, when reconfigurations are infrequent, the performance of the implementation is not significantly affected. Since the main aim of our algorithm is to reduce the size of the gossip messages, we decided not to reconfigure during the experiment. Future experiments will include reconfiguration activities; however, it is also important to evaluate the system when the reconfiguration traffic does not interfere with the routine gossip.

Second, the experiments were conducted using our preliminary implementation. This implementation is a proof-of-concept for the methodology developed in [3] to translate the IOA specification into Java code. Both RAMBO and LL-RAMBO have been implemented faithfully to their specifications; however, no attempt was made to optimize either system. Therefore, the performance improvements in Fig. 4(a) and (b) should be considered in relative terms.

Finally, observe that due to node failures, some gossip messages may continue to grow without bound (since the sender does not receive acknowledgments from the failed node). To diminish the communication and processing impact in this case, we use *exponential backoff*: if a node does not receive a gossip message from some other node, it will double the delay between consecutive gossip rounds to that node; normal timing is restored once communication is reestablished. Note that both algorithms may take advantage of this assumption. In our experiments, gossip messages were not sent to the virtual-failed nodes (this approximates exponential backoff in the case of permanently crashed nodes). The results in Fig. 4(a) and (b) illustrate the performance of RAMBO and LL-RAMBO when gossip messages are exchanged only between non-failed nodes.

8. Discussion and future work

We presented an algorithm for long-lived atomic data in dynamic networks. Prior solutions for dynamic networks [1,2] did not allow the participants to leave gracefully, and relied on gossip that involved sending messages whose size grew with time. The new algorithm, called LL-RAMBO improves on prior work by supporting the graceful departures of participants and by implementing incremental gossip. The algorithm substantially reduces the size and the number of gossip messages, leading to a much improved performance of the read and write operations. We show that the improved algorithm implements atomic objects in the presence of arbitrary asynchrony, dynamic node joins, and departures. The algorithm relies on simple point-to-point channels that permit message loss and reordering.

We analyzed the efficiency of the algorithm and illustrated its performance using our implementation of the algorithm in the local-area setting. In trading knowledge for communication, the algorithm increases the local memory requirements. Specifically, the needed local storage is quadratic in the number of active participants. Future plans include exploring optimizations that will reduce the local storage usage to be about linear in the number of participants. Also, we plan to further explore ways to decrease the number of gossip messages sent. Currently, LL-RAMBO allows all-to-all exchange of information among active participants, and when gossip is unconstrained and communication bandwidth is limited, network congestion may degrade the performance of such a system. One way to solve this problem is to constrain the gossip patterns. Restricting gossip vacuously preserves the linearizability property of the algorithm, but will possibly alter its performance.

Acknowledgements

The authors thank Nancy Lynch and Seth Gilbert for many discussions. The authors are also grateful to the anonymous reviewers whose comments helped improve the presentation of the paper.

Appendix. RAMBO specification

See Figs. A.1–A.8.

Signature:

Input:

$\text{join}(\text{rambo}, J)_i$, J a finite subset of $I - \{i\}$
 $\text{join-ack}(r)_i$, $r \in \{\text{recon}, \text{rw}\}$
 fail_i

State:

$\text{status} \in \{\text{idle}, \text{joining}, \text{active}\}$, initially *idle*
 $\text{child-status} \in \{\text{recon}, \text{rw}\} \rightarrow \{\text{idle}, \text{joining}, \text{active}\}$, initially everywhere *idle*
 $\text{hints} \subseteq I$, initially \emptyset
 failed , a Boolean, initially *false*

Transitions:

Input $\text{join}(\text{rambo}, J)_i$

Effect:

if $\neg \text{failed}$ then
 if $\text{status} = \text{idle}$ then
 $\text{status} \leftarrow \text{joining}$
 $\text{hints} \leftarrow J$

Output $\text{send}(\text{join})_{i,j}$

Precondition:

$\neg \text{failed}$
 $\text{status} = \text{joining}$
 $j \in \text{hints}$

Effect:

none

Output $\text{join}(r)_i$

Precondition:

$\neg \text{failed}$
 $\text{status} = \text{joining}$
 $\text{child-status}(r) = \text{idle}$

Effect:

$\text{child-status}(r) \leftarrow \text{joining}$

Input $\text{join-ack}(r)_i$

Effect:

if $\neg \text{failed}$ then
 if $\text{status} = \text{joining}$ then
 $\text{child-status}(r) \leftarrow \text{active}$

Output:

$\text{send}(\text{join})_{i,j}$, $j \in I - \{i\}$
 $\text{join}(r)_i$, $r \in \{\text{recon}, \text{rw}\}$
 $\text{join-ack}(\text{rambo})_i$

Output $\text{join-ack}(\text{rambo})_i$

Precondition:

$\neg \text{failed}$
 $\text{status} = \text{joining}$
 $\forall r \in \{\text{recon}, \text{rw}\} : \text{child-status}(r) = \text{active}$

Effect:

$\text{status} \leftarrow \text{active}$

Input fail_i

Effect:

$\text{failed} \leftarrow \text{true}$

Fig. A.1. *Joiner_i*: Signature, state, and transitions.

Signature:

Input:

read_i
 $\text{write}(v)_i$, $v \in V$
 $\text{new-config}(c, k)_i$, $c \in C$, $k \in \mathbb{N}^+$
 $\text{recv}(\text{join})_{j,i}$, $j \in I - \{i\}$
 $\text{recv}(m)_{j,i}$, $m \in M$, $j \in I$
 $\text{join}(\text{rw})_i$
 fail_i

State:

$\text{status} \in \{\text{idle}, \text{joining}, \text{active}\}$, initially *idle*
 world , a finite subset of I , initially \emptyset
 $\text{value} \in V$, initially v_0
 $\text{tag} \in T$, initially $(0, i_0)$
 $\text{cmap} \in \text{CMap}$, initially $\text{cmap}(0) = c_0$,
 $\text{cmap}(k) = \perp$ for $k \geq 1$
 $\text{pnum1} \in \mathbb{N}$, initially 0
 $\text{pnum2} \in I \rightarrow \mathbb{N}$, initially everywhere 0
 failed , a Boolean, initially *false*

Output:

$\text{join-ack}(\text{rw})_i$
 $\text{read-ack}(v)_i$, $v \in V$
 write-ack_i
 $\text{send}(m)_{i,j}$, $m \in M$, $j \in I$

Internal:

query-fix_i
 prop-fix_i
 $\text{cfg-upgrade}(k)_i$, $k \in \mathbb{N}^{>0}$
 $\text{cfg-upg-query-fix}(k)_i$, $k \in \mathbb{N}^{>0}$
 $\text{cfg-upg-prop-fix}(k)_i$, $k \in \mathbb{N}^{>0}$
 $\text{cfg-upgrade-ack}(k)_i$, $k \in \mathbb{N}^{>0}$

op , a record with fields:

$\text{type} \in \{\text{read}, \text{write}\}$
 $\text{phase} \in \{\text{idle}, \text{query}, \text{prop}, \text{done}\}$, initially *idle*
 $\text{pnum} \in \mathbb{N}$
 $\text{cmap} \in \text{CMap}$
 acc , a finite subset of I
 $\text{value} \in V$

upg , a record with fields:

$\text{phase} \in \{\text{idle}, \text{query}, \text{prop}\}$, initially *idle*
 $\text{pnum} \in \mathbb{N}$
 $\text{cmap} \in \text{CMap}$,
 acc , a finite subset of I
 $\text{target} \in \mathbb{N}$

Fig. A.2. *Reader–Writer_i*: Signature and state.

Input $\text{join}(\text{rw})_i$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{idle}$ then if $i = i_0$ then $\text{status} \leftarrow \text{active}$ else $\text{status} \leftarrow \text{joining}$ $\text{world} \leftarrow \text{world} \cup \{i\}$	Input $\text{recv}(\text{join})_{j,i}$ Effect: if $\neg\text{failed}$ then if $\text{status} \neq \text{idle}$ then $\text{world} \leftarrow \text{world} \cup \{j\}$	Output $\text{join-ack}(\text{rw})_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ Effect: none
--	---	---

Fig. A.3. *Reader-Writer_i*: Join-related transitions.

Input $\text{cfg-upgrade}(k)_i$ Effect: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{upg.phase} = \text{idle}$ $\text{cmap}(k) \in C$ $\forall \ell \in \mathbb{N}, \ell < k : \text{cmap}(\ell) \neq \perp$ $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\text{upg} \leftarrow \langle \text{query}, \text{pnum1}, \text{cmap}, \emptyset, k \rangle$ Internal $\text{cfg-upg-query-fix}(k)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{upg.phase} = \text{query}$ $\text{upg.target} = k$ $\forall \ell \in \mathbb{N}, \ell < k : \text{upg.cmap}(\ell) \in C$ $\Rightarrow \exists R \in \text{read-quorums}(\text{upg.cmap}(\ell)) :$ $\exists W \in \text{write-quorums}(\text{upg.cmap}(\ell)) :$ $R \cup W \subseteq \text{upg.acc}$ Effect: $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\text{upg.pnum} \leftarrow \text{pnum1}$ $\text{upg.phase} \leftarrow \text{transient}$ $\text{upg.acc} \leftarrow \emptyset$	Internal $\text{cfg-upg-prop-fix}(k)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{upg.phase} = \text{prop}$ $\text{upg.target} = k$ $\exists W \in \text{write-quorums}(\text{upg.cmap}(k)) : W \subseteq \text{upg.acc}$ Effect: for $\ell \in \mathbb{N} : \ell < k$ do $\text{cmap}(\ell) \leftarrow \perp$ Internal $\text{cfg-upgrade-ack}(k)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{upg.target} = k$ $\forall \ell \in \mathbb{N}, \ell < k : \text{cmap}(\ell) = \perp$ Effect: $\text{upg.phase} \leftarrow \text{idle}$
--	---

Fig. A.4. *Reader-Writer_i*: Configuration-management transitions.

Output $\text{send}((W, v, t, \text{cm}, \text{pns}, \text{pnr}))_{i,j}$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $j \in \text{world}$ $\langle W, v, t, \text{cm}, \text{pns}, \text{pnr} \rangle =$ $\langle \text{world}, \text{value}, \text{tag}, \text{cmap}, \text{pnum1}, \text{pnum2}(j) \rangle$ Effect: none Input $\text{recv}((W, v, t, \text{cm}, \text{pns}, \text{pnr}, m))_{j,i}$ Effect: if $\neg\text{failed}$ then if $\text{status} \neq \text{idle}$ then $\text{status} \leftarrow \text{active}$ $\text{world} \leftarrow \text{world} \cup W$ if $t > \text{tag}$ then $\langle \text{value}, \text{tag} \rangle \leftarrow \langle v, t \rangle$ $\text{cmap} \leftarrow \text{update}(\text{cmap}, \text{cm})$ $\text{pnum2}(j) \leftarrow \max(\text{pnum2}(j), \text{pns})$ if $\text{op.phase} \in \{\text{query}, \text{prop}\} \wedge \text{pnr} \geq \text{op.pnum}$ then $\text{op.cmap} \leftarrow \text{extend}(\text{op.cmap}, \text{truncate}(\text{cm}))$ if $\text{op.cmap} \in \text{Truncated}$ then $\text{op.acc} \leftarrow \text{op.acc} \cup \{j\}$ else $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\text{op.acc} \leftarrow \emptyset$ $\text{op.cmap} \leftarrow \text{truncate}(\text{cmap})$ if $\text{upg.phase} \neq \text{idle}$ and $\text{pnr} \geq \text{upg.pnum}$ then $\text{upg.acc} \leftarrow \text{upg.acc} \cup \{j\}$	Internal query-fix_i Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{op.type} \in \{\text{read}, \text{write}\}$ $\text{op.phase} = \text{query}$ $\forall k \in \mathbb{N}, c \in C : (\text{op.cmap}(k) = c)$ $\Rightarrow (\exists R \in \text{read-quorums}(c) :$ $R \subseteq \text{op.acc})$ Effect: if $\text{op.type} = \text{read}$ then $\text{op.value} \leftarrow \text{value}$ else $\text{value} \leftarrow \text{op.value}$ $\text{tag} \leftarrow \langle \text{tag.seq} + 1, i \rangle$ $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\text{op.pnum} \leftarrow \text{pnum1}$ $\text{op.phase} \leftarrow \text{prop}$ $\text{op.cmap} \leftarrow \text{truncate}(\text{cmap})$ $\text{op.acc} \leftarrow \emptyset$ Internal prop-fix_i Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{op.type} \in \{\text{read}, \text{write}\}$ $\text{op.phase} = \text{prop}$ $\forall k \in \mathbb{N}, c \in C : (\text{op.cmap}(k) = c)$ $\Rightarrow (\exists W \in \text{write-quorums}(c) :$ $W \subseteq \text{op.acc})$
--	---

Fig. A.5. *Reader-Writer_i*: Read/write and failure transitions.

<p>Input $\text{new-config}(c, k)_i$ Effect: if $\neg \text{failed}$ then if $\text{status} \neq \text{idle}$ then $\text{cmap}(k) \leftarrow \text{update}(\text{cmap}(k), c)$</p> <p>Input read_i Effect: if $\neg \text{failed}$ then if $\text{status} \neq \text{idle}$ then $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\langle \text{op.pnum}, \text{op.type}, \text{op.phase} \rangle$ $\leftarrow \langle \text{pnum1}, \text{read}, \text{query} \rangle$ $\langle \text{op.cmap}, \text{op.acc} \rangle$ $\leftarrow \langle \text{truncate}(\text{cmap}), \emptyset \rangle$</p> <p>Input $\text{write}(v)_i$ Effect: if $\neg \text{failed}$ then if $\text{status} \neq \text{idle}$ then $\text{pnum1} \leftarrow \text{pnum1} + 1$ $\langle \text{op.pnum}, \text{op.type}, \text{op.phase} \rangle$ $\leftarrow \langle \text{pnum1}, \text{write}, \text{query} \rangle$ $\langle \text{op.cmap}, \text{op.acc}, \text{op.value} \rangle$ $\leftarrow \langle \text{truncate}(\text{cmap}), \emptyset, v \rangle$</p>	<p>Effect: $\text{op.phase} = \text{done}$</p> <p>Output $\text{read-ack}(v)_i$ Precondition: $\neg \text{failed}$ $\text{status} = \text{active}$ $\text{op.type} = \text{read}$ $\text{op.phase} = \text{done}$ $v = \text{op.value}$ Effect: $\text{op.phase} = \text{idle}$</p> <p>Output write-ack_i Precondition: $\neg \text{failed}$ $\text{status} = \text{active}$ $\text{op.type} = \text{write}$ $\text{op.phase} = \text{done}$ Effect: $\text{op.phase} = \text{idle}$</p> <p>Input fail_i Effect: $\text{failed} \leftarrow \text{true}$</p>
--	--

Fig. A.5. Continued

<p>Input: $\text{init}(v)_{k,c,i}, v \in V, i \in \text{members}(c)$ $\text{fail}_i, i \in \text{members}(c)$</p>	<p>Output: $\text{decide}(v)_{k,c,i}, v \in V, i \in \text{members}(c)$</p>
--	--

Fig. A.6. $\text{Cons}(k, c)$: External signature (an external consensus service).**Signature:**

<p>Input: $\text{join}(\text{recon})_i$ $\text{recon}(c, c')_{k,i}, c, c' \in C, i \in \text{members}(c)$ $\text{decide}(c)_{k,i}, c \in C, k \in \mathbb{N}^+$ $\text{recv}(\langle \text{config}, c, k \rangle)_{j,i}, c \in C, k \in \mathbb{N}^+,$ $i \in \text{members}(c), j \in I - \{i\}$ $\text{recv}(\langle \text{init}, c, c', k \rangle)_{j,i}, c, c' \in C, k \in \mathbb{N}^+,$ $i, j \in \text{members}(c), j \neq i$ fail_i</p>	<p>Output: $\text{join-ack}(\text{recon})_i$ $\text{new-config}(c, k)_i, c \in C, k \in \mathbb{N}^+$ $\text{init}(c, c')_{k,i}, c, c' \in C, k \in \mathbb{N}^+, i \in \text{members}(c)$ $\text{recon-ack}(b)_i, b \in \{\text{ok}, \text{nok}\}$ $\text{report}(c)_i, c \in C$ $\text{send}(\langle \text{config}, c, k \rangle)_{i,j}, c \in C, k \in \mathbb{N}^+,$ $j \in \text{members}(c) - \{i\}$ $\text{send}(\langle \text{init}, c, c', k \rangle)_{i,j}, c, c' \in C, k \in \mathbb{N}^+,$ $i, j \in \text{members}(c), j \neq i$</p>
--	---

State:

<p>$\text{status} \in \{\text{idle}, \text{active}\}$, initially idle. $\text{rec-cmap} \in \text{CMap}$, initially $\text{rec-cmap}(0) = c_0$ and $\text{rec-cmap}(k) = \perp$ for all $k \neq 0$. $\text{did-new-config} \subseteq \mathbb{N}^+$, initially \emptyset $\text{reported} \subseteq C$, initially \emptyset</p>	<p>$\text{op-status} \in \{\text{idle}, \text{active}\}$, initially idle $\text{op-outcome} \in \{\text{ok}, \text{nok}, \perp\}$, initially \perp $\text{cons-data} \in (\mathbb{N}^+ \rightarrow (C \times C))$, initially everywhere \perp $\text{did-init} \subseteq \mathbb{N}^+$, initially \emptyset failed, a Boolean, initially false</p>
---	---

Fig. A.7. Recon_i : Signature and state.

<p>Input $\text{join}(\text{recon})_i$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{idle}$ then $\text{status} \leftarrow \text{active}$</p> <p>Output $\text{join-ack}(\text{recon})_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$</p> <p>Effect: none</p> <p>Output $\text{new-config}(c, k)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{rec-cmap}(k) = c$ $k \notin \text{did-new-config}$</p> <p>Effect: $\text{did-new-config} \leftarrow \text{did-new-config} \cup \{k\}$</p> <p>Output $\text{send}((\text{config}, c, k))_{i,j}$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{rec-cmap}(k) = c$</p> <p>Effect: none</p> <p>Input $\text{recv}((\text{config}, c, k))_{j,i}$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{active}$ then $\text{rec-cmap}(k) \leftarrow c$</p> <p>Output $\text{report}(c)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $c = \text{rec-cmap}(k)$ $\forall \ell > k : \text{rec-cmap}(\ell) = \perp$ $c \notin \text{reported}$</p> <p>Effect: $\text{reported} \leftarrow \text{reported} \cup \{c\}$</p> <p>Input $\text{recon}(c, c')_i$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{active}$ then $\text{op-status} \leftarrow \text{active}$ let $k = \max(\{\ell : \text{rec-cmap}(\ell) \in C\})$ if $c = \text{rec-cmap}(k) \wedge \text{cons-data}(k+1) = \perp$ then $\text{cons-data}(k+1) \leftarrow (c, c')$ $\text{op-outcome} \leftarrow \perp$ else $\text{op-outcome} \leftarrow \text{nok}$</p>	<p>Output $\text{init}(c')_{k,c,i}$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{cons-data}(k) = (c, c')$ if $k \geq 1$ then $k-1 \in \text{did-new-config}$ $k \notin \text{did-init}$</p> <p>Effect: $\text{did-init} \leftarrow \text{did-init} \cup \{k\}$</p> <p>Output $\text{send}((\text{init}, c, c', k))_{i,j}$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{cons-data}(k) = (c, c')$ $k \in \text{did-init}$</p> <p>Effect: none</p> <p>Input $\text{recv}((\text{init}, c, c', k))_{j,i}$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{active}$ then if $\text{rec-cmap}(k-1) = \perp$ then $\text{rec-cmap}(k-1) \leftarrow c$ if $\text{cons-data}(k) = \perp$ then $\text{cons-data}(k) \leftarrow (c, c')$</p> <p>Input $\text{decide}(c')_{k,c,i}$ Effect: if $\neg\text{failed}$ then if $\text{status} = \text{active}$ then $\text{rec-cmap}(k) \leftarrow c'$ if $\text{op-status} = \text{active}$ then if $\text{cons-data}(k) = (c, c')$ then $\text{op-outcome} \leftarrow \text{ok}$ else $\text{op-outcome} \leftarrow \text{nok}$</p> <p>Output $\text{recon-ack}(b)_i$ Precondition: $\neg\text{failed}$ $\text{status} = \text{active}$ $\text{op-status} = \text{active}$ $\text{op-outcome} = b$</p> <p>Effect: $\text{op-status} = \text{idle}$</p> <p>Input fail_i Effect: $\text{failed} \leftarrow \text{true}$</p>
--	--

Fig. A.8. Recon_i : Transitions.

References

- [1] N. Lynch, A.A. Shvartsman, RAMBO: A reconfigurable atomic memory service for dynamic networks, in: Proc. of 16th International Symposium on Distributed Computing, 2002, pp. 173–190.
- [2] S. Gilbert, N. Lynch, A.A. Shvartsman, RAMBO II: Rapidly reconfigurable atomic memory for dynamic networks, in: Proc. of International Conference on Dependable Systems and Networks, 2003, pp. 259–268.
- [3] P.M. Musial, A.A. Shvartsman, Implementing a reconfigurable atomic memory service for dynamic networks, in: Proc. of 18'th International Parallel and Distributed Symposium — FTPDS WS, 2004, p. 208b.
- [4] N.A. Lynch, M. Tuttle, Hierarchical correctness proofs for distributed algorithms, Technical Report, 1987.
- [5] D.K. Gifford, Weighted voting for replicated data, in: Proc. of 7th ACM Symp. on Oper. Sys. Princ., 1979, pp. 150–162.
- [6] R. Thomas, A majority consensus approach to concurrency control for multiple copy databases, ACM Transactions on Database Systems 4 (2) (1979) 180–209.

- [7] E. Upfal, A. Wigderson, How to share memory in a distributed system, *Journal of the ACM* 34 (1) (1987) 116–127.
- [8] B. Awerbuch, P. Vitanyi, Atomic shared register access by asynchronous hardware, in: *Proc. of 27th IEEE Symposium on Foundations of Computer Science*, 1986, pp. 233–243.
- [9] H. Attiya, A. Bar-Noy, D. Dolev, Sharing memory robustly in message passing systems, *Journal of the ACM* 42 (1) (1996) 124–142.
- [10] B. Englert, A.A. Shvartsman, Graceful quorum reconfiguration in a robust emulation of shared memory, in: *Proc. of International Conference on Distributed Computer Systems*, 2000, pp. 454–463.
- [11] N. Lynch, A.A. Shvartsman, Robust emulation of shared memory using dynamic quorum-acknowledged broadcasts, in: *Proc. of 27th Int'l Symp. on Fault-Tolerant Comp.*, 1997, pp. 272–281.
- [12] K. Birman, T. Joseph, Exploiting virtual synchrony in distributed systems, in: *Proc. of the 11th ACM Symposium on Operating Systems Principles*, December 1987.
- [13] On group communication services, *Communications of the ACM* 39 (4) (1996) (special issue).
- [14] A. El Abbadi, F. Cristian, D. Skeen, An efficient fault-tolerant protocol for replicated data management, in: *Proc. of 4th ACM Symposium on Principles of Databases*, 1985, pp. 215–228.
- [15] R. De Prisco, A. Fekete, N. Lynch, A.A. Shvartsman, A dynamic primary configuration group communication service, in: *Proc. of 13th Int'l Conference on Distributed Computing*, 1999, pp. 64–78.
- [16] S. Jajodia, D. Mutchler, Dynamic voting algorithms for maintaining the consistency of a replicated database, *Transactions on Database Systems* 15 (2) (1990) 230–280.
- [17] D. Dolev, E. Lotem, I. Keidar, Dynamic voting for consistent primary components, in: *Proc. of 16th ACM Symposium on Principles of Distributed Computing*, 1997, pp. 63–71.
- [18] M.P. Herlihy, Dynamic quorum adjustment for partitioned data, *ACM Transactions on Database Systems* 12 (2) (1987) 170–194.
- [19] K.P. Birman, M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu, Y. Minsky, Bimodal multicast, *ACM Transactions on Computer Systems* 17 (2) (1999) 41–88.
- [20] O.M. Cheiner, A.A. Shvartsman, Implementing an eventually-serializable data service as a distributed system building block, vol. 45, 1999, pp. 43–71.
- [21] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, D. Terry, Epidemic algorithms for replicated database maintenance, in: *Proc. of 6th ACM Symp. on Princ. of Distr. Comp.*, 1987, pp. 1–12.
- [22] A. Ghorbani, V. Bhavsar, Training artificial neural networks using variable precision incremental communication, in: *Proc. of IEEE World Congress On Computational Intelligence*, vol. 3, 1994, pp. 1409–1414.
- [23] K. Petersen, M.J. Spreitzer, D.B. Terry, M.M. Theimer, Bayou: Replicated database services for world-wide applications, in: *Proc. of 7th ACM SIGOPS European Workshop*, 1996, pp. 275–280.
- [24] Y. Minsky, Spreading rumors cheaply, quickly, and reliably, Ph.D. Thesis, Cornell University, 2002.
- [25] R. Van Renesse, K.P. Birman, W. Vogels, Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining, *ACM Transactions on Computer Systems* 21 (2) (2003) 164–206.
- [26] R.A. Golding, Weak-consistency group communication and membership, Ph.D. Thesis, University of California, 1992.
- [27] M. Rabinovich, N. Gehani, A. Kononov, Efficient update propagation in epidemic replicated databases, in: *Proc. of 5th Int. Conf. on Extending Database Technology*, 1996, pp. 207–222.
- [28] R.G. Guy, J.S. Heidemann, W. Mak, T.W. Page Jr., G.J. Popek, D. Rothmeier, Implementation of the ficus replicated file system, in: *Proc. of Summer USENIX Conference*, 1990, pp. 63–71.
- [29] C. Georgiou, P.M. Musial, A.A. Shvartsman, Long-lived RAMBO: Trading knowledge for communication, in: *Proc. of 11'th Colloquium on Structural Information and Communication Complexity*, Springer, 2004, pp. 185–196.
- [30] N.A. Lynch, *Distributed Algorithms*, Morgan Kaufmann Publishers, 1996.
- [31] L. Lamport, The part-time parliament, *ACM Transactions on Computer Systems* 16 (2) (1998) 133–169.
- [32] Sh. Dolev, S. Gilbert, N. Lynch, A. Shvartsman, J. Welch, Geoquorums: Implementing atomic memory in ad hoc networks, in: *Proc. of 17th International Symposium on Distributed Computing*, 2003, pp. 306–320.