SCALABLE COMPUTING: PRACTICE AND EXPERIENCE
Volume 7, Number 4, pp. 53–87. http://www.scpe.org

ISSN 1895-1767
© 2006 SWPS

A FAULT-TOLERANT DIRECTORY SERVICE FOR MOBILE AGENTS BASED ON FORWARDING POINTERS

LUC MOREAU*

Key words. mobile agents, distributed directory service, fault tolerance

Abstract. A reliable communication layer is an essential component of a mobile agent system. We present a new fault-tolerant directory service for mobile agents, which can be used to route messages reliably to them, even in the presence of failures of intermediary nodes between senders and receivers. The directory service, based on a technique of forwarding pointers, introduces some redundancy in order to ensure resilience to stopping failures of nodes containing forwarding pointers; in addition, it avoids cyclic routing of messages, and it supports a technique to collapse chains of pointers that allows direct communication between agents. We have formalised the algorithm and derived a *fully mechanical proof* of its correctness using the proof assistant Coq; we report on our experience of designing the algorithm and deriving its proof of correctness. The complete source code of the proof is made available from the WWW.

1. Introduction. While mobile agents have been touted as a major programming paradigm for structuring distributed applications [3, 5], several important issues remain to be addressed before mobile agents can become a mainstream technology for such applications: among them, a *communication system* and a *security infrastructure* are needed respectively for facilitating communications between mobile agents and for protecting agents and their hosts.

In this article, we focus solely on the problem of communications, whereas security issues are the focus of other publications, such as [20, 21]. Various authors have previously investigated a communication layer for mobile agents based on *forwarding pointers* [17, 10]. In such an approach, when mobile agents migrate, they leave forwarding pointers that are used to route messages. This simple approach raises some concerns: first, one needs to avoid cyclic routing when agents migrate to previously visited sites; second, chains of pointers can become arbitrarily long and increase the cost of communication. The first problem can be addressed by using timestamps [10], whereas the second can be solved by techniques such as lazy updates and piggy-backing of information on messages [14]. For structuring and clarity purposes, a communication layer for mobile agents is usually defined in terms of a message router and a directory service; the latter tracks mobile agents' locations, whereas the former forwards messages using the information provided by the latter.

Directory services based on forwarding pointers are currently *not tolerant* to failures [17, 10]: the failure of a node containing a forwarding pointer may prevent finding agents' positions. The purpose of this article is to present a directory service, fully distributed and resilient to failures exhibited by intermediary nodes that possibly contain forwarding pointers. This algorithm may be used to specify fault-tolerant message routers.

We consider stopping failures according to which processes are allowed to stop during the course or their execution [7]. The essence of our fault-tolerant distributed directory service is to introduce *redundancy* of forwarding pointers, typically by making N copies of agents' location information. This type of redundancy ensures the resilience of the algorithm to a maximum of $N-1$ failures of intermediary nodes. We show that the complexity of the algorithm remains linear in N. Our specific contributions are:

1. A *new directory service* based on forwarding pointers that is fault-tolerant, preventing cyclic routing, and not involving any static location;
2. A *full mechanical proof* of its correctness, using the proof assistant Coq [1]; the complete source code of the proof (involving some 25000 tactic invocations) may be downloaded from the following URL [9].

This article is an extended version of a paper originally published at the AIMS track (Agents, interactions, mobility, and systems) at the ACM Symposium on Applied Computing 2002 [11]. It extends the original paper with a series of graphical animations that illustrate the behaviour of the algorithm.

We begin this paper by a survey of background work (Section 2) and follow by a summary of a routing algorithm based on forwarding pointers (Section 3). We present our new directory service (Section 4) and its formalisation as an abstract machine (Section 5). The purpose of Section 6 is to summarise the correctness properties of the algorithm: its safety states that the distributed directory service correctly and uniquely identifies agents' positions, whereas the liveness property shows that the algorithm reaches a stable state after a finite number of transitions, once agents stop migrating. Then, in Section 7, we report on our experience of designing the algorithm and deriving its proof of correctness, and we suggest possible variants or extensions, before discussing further related work (Section 8) and concluding the paper.

*School of Electronics and Computer Science,University of Southampton, Southampton SO17 1BJ UK, L.Moreau@ecs.soton.ac.uk

2. Background. The topic of mobile agent tracking and communication has been researched extensively by the mobile agent community. Very early on, location-aware communications were proposed: they consist of sending messages to locations where agents are believed to be, but typically result in failure if the receiver agent has migrated [16, 22].

For a number of applications, such a kind of communication layer is not satisfactory when the expectation is to deliver messages *reliably* to a recipient, wherever its location and whatever the route adopted; for instance, it would not be acceptable to lose a message in a negotiation between two mobile agents, simply because one of them changed location. To combat this problem, location-transparent communication services were introduced as a means to route and deliver messages automatically to mobile agents, independently of their migration. Such services have been shown to be implementable on top of a location-aware communication layer [22].

In the category of location-transparent communication layers, there are essentially two distinct approaches, respectively based on *home agents* and *forwarding pointers*. In systems based on home agents, such as Aglets [5], each mobile agent is associated with a non-mobile home agent. In order to communicate with a mobile agent, a message has to be sent to its associated home agent, which forwards it to the mobile one; when a mobile agent migrates, it informs its home agent of its new position. Alternatively, in mobile agent systems such as Voyager [17], agents that migrate leave trails of forwarding pointers, which are used to route messages.

In pervasive computing environments, for example, the mechanism of a home agent may defeat the purpose of using mobile agents by re-introducing centralisation: the home agent approach puts a burden on the infrastructure, which may hamper its scalability, in particular, in massively distributed systems. A typical illustration of the difficulty consists of two mobile agents with respective home bases in the US and Europe having to communicate with each other, while located at a host in Australia. In such a scenario, routing via home agents is not desirable, and may not be possible when the Australian host is temporarily disconnected from the network. If we introduce a mechanism by which home agents change location dynamically according to the task at hand, we face the problem of how to communicate reliably with a home agent, which is itself mobile. Alternatively, we could only use the home agent to bootstrap communication, and then shortcut the route, but this approach becomes unreliable once agents migrate. Finally, the home agent is also as a *single point of failure*: when it exhibits a failure, it becomes impossible to track the mobile agent or to route messages to it.

A naive forwarding pointer implementation causes communications to become more expensive as agents migrate, because chains of pointers increase. Chains of pointers need to be collapsed promptly so that mobile agents become independent of the hosts they previously visited. Once the chain has collapsed direct communications become possible and avoid the awkward scenario discussed above. As far as tolerance to failures is concerned, the failure of an intermediary node with a forwarding pointer prevents upstream nodes from forwarding messages. Hence, collapsing chains of pointers is crucial to reduce the system's exposure to failures.

Coordination models offer a more asynchronous form of communication, typically involving a tuple space [4]. As tuple spaces are non-mobile, they may suffer from the same problem as the home agent. To improve locality, tuple spaces may be distributed, and coordinated by replication protocols, but maintaining consistency in the presence of updates is a non-trivial problem. A further inconvenience of the coordination approach is that it requires coordinated processes to poll tuple spaces, which may be an inefficient approach in terms of both communication and computation. To combat this problem, tuple spaces generally provide a mechanism by which registered clients can be notified of the arrival of a new tuple: when clients are mobile, we are back to the problem of delivering notifications reliably to a potentially mobile recipient. Alternatively, if the tuple space itself is mobile [18], the problem is then to deliver messages to the tuple space.

This discussion shows that reliable delivery of messages to mobile agents without using static locations to route messages is essential, even if peer-to-peer communications are not adopted as the high-level interaction paradigm between agents. Previous work has focused on formalisation [10] and implementation [17] of forwarding pointers, but such solutions were not fault-tolerant. We summarise such an approach in Section 3 before extending it with support for failures in Section 5.

3. Original Algorithm. In this section, we summarise the principles of a communication layer based on forwarding pointers [10] without any fault-tolerance. The algorithm comprises two components: a distributed directory service and a message router, which we describe below; we then describe why the algorithm is not tolerant to failures.

3.1. Distributed Directory Service. Each mobile agent is associated with a timestamp that is increased every time the agent migrates. When an agent has decided to migrate to a new location, it requests the communication layer to transport it to its new destination. When the agent arrives at a new location, an acknowledgement message containing both its new position and its newly-incremented timestamp is sent to its previous location. As a result, for each site, one of the following three cases is valid for each agent A: *(i)* the agent A is local, *(ii)* the agent A is in transit but has not acknowledged its new position yet, or *(iii)* the agent A is known to have been at a remote location with a given timestamp.

Figures 3.1 to 3.8 present an animation of the algorithm by showing various agent locations and associated site states. (*Note: such animation is best viewed in colour directly in a pdf viewer.*) Each site contains a local state composed of four components, which we intuitively introduce now.

1. *mob*: the lastest known mobility counter, i. e. timestamp, of the agent;
2. *loc*: the lastest known location of the agent;
3. *pres*: a boolean flag indicating if the agent is present locally or not;
4. *ack*: whether an acknowledgement message has to be sent or not.

As an illustration, in Figure 3.1, at site s_2, agent is known to have timestamp $t+1$ and location s_2, and to be present locally. No acknowledgement needs to be sent by s_2.

The mobile agent located at site s_2 decides to migrate to a new location (Figure 3.2). As it makes the request to the transport layer to be transported and leaves s_2 (Figure 3.3), the agent state is packaged up with its previous location's name s_2 and the timestamp it would have at its next location $t+2$. The latest known timestamp and location remain unchanged at s_2 since the agent's presence elsewhere has not been confirmed yet. However, the agent is now known to be absent from s_2.

As the agent arrives at s_3 (Figure 3.4), s_3 acquires the knowledge that the agent is present locally (hence, reflected in the states *mob*, *loc* and *pres*). Furthermore, the state *ack* is changed at s_3 to indicate that the new agent's position has to be communicated to its previous location s_2.

The sending of the acknowledgement message by s_3 (Figure 3.5) clears the *ack* state at s_3 and results in the latest known location of the agent to be communicated to s_2 (as reflected by the change of *mob* and *loc* at s_2 in Figure 3.6). As a result, the processing of the acknowledgement message by s_2 results in a *forwarding pointer* being set up from s_2 to s_3.

Timestamps are essential to avoid race conditions between acknowledgement messages: by using timestamps, a site can decide which position information is the most recent, and therefore can avoid creating cycles in the graph of forwarding pointers (see [10] for details).

In order to avoid an increasing cost of communication when the agent migrates, it is useful to reduce the length of forwarding pointers. To this end, we have established [10] that information messages, containing a site's belief about the agent's position and its timestamp, can be communicated by any site to any other site. The processing of such inform messages, like acknowledgement messages, updates the receipient's local state if the knowledge received is more recent than the local one (as per indicated by the timestamp). Figure 3.7 illustrates the sending of such an information message by s_2, sharing its knowledge of the agent's position with s_1. On receipt of the information message (Figure 3.8), s_1 updates its internal tables. By propagating such agent's positions, one can reduces the length of chains of pointers. Different strategies such as eager or lazy propagation are discussed in [14].

3.2. Message Router. Sites rely on the information about agents' positions in order to route messages. For any incoming message aimed at an agent A, the message will be delivered to A if A is known to be local. If A is in transit, the message will be enqueued, until A's location becomes known; otherwise, the message is forwarded to A's latest known location.

3.3. Absence of Fault Tolerance. There is no redundancy in the information concerning an agent's location. Indeed, sites only remember the most recent location of an agent, and only the previous agent's location is informed of the new agent's position after a migration. As a result, a site (transitively) pointing at a site exhibiting a failure has lost its route to the agent.

mob(s1)= t+1
loc(s1)= s2
pres(s1)=false
ack(s1)=negative

mob(s2)=t+1
loc(s2)=s2
pres(s2)=true
ack(s2)=negative

mob(s3)=_
loc(s3)=_
pres(s3)=false
ack(s3)=negative

s1

s2

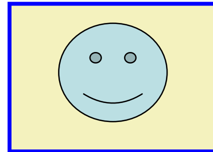

s3

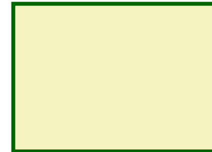

FIG. 3.1. *Original Algorithm (1/8)*

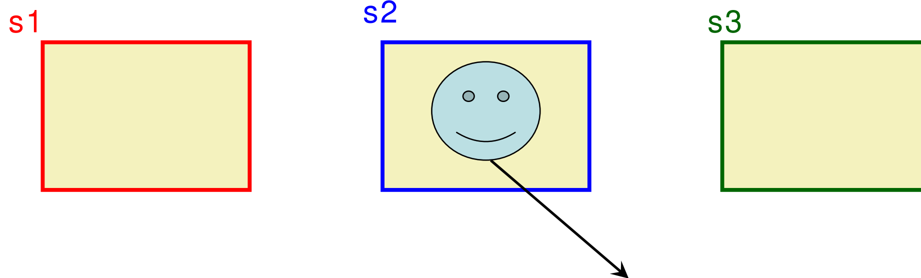

FIG. 3.2. *Original Algorithm (2/8)*

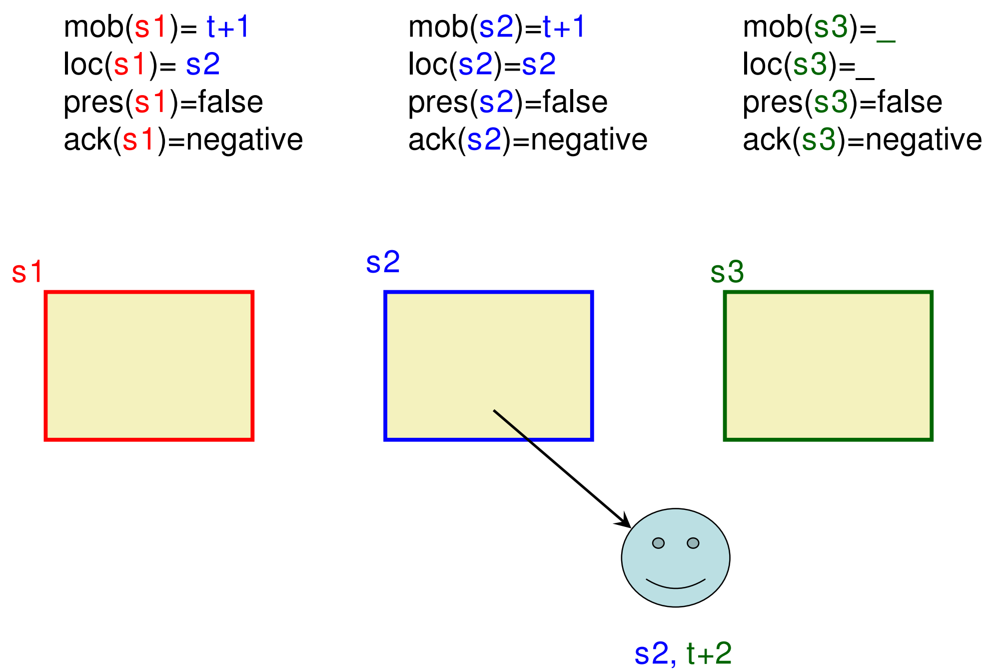

FIG. 3.3. *Original Algorithm (3/8)*

mob(s1)= t+1
loc(s1)= s2
pres(s1)=false
ack(s1)=negative

mob(s2)=t+1
loc(s2)=s2
pres(s2)=false
ack(s2)=negative

mob(s3)=t+2
loc(s3)=s3
pres(s3)=true
ack(s3)=positive(s2, t+2)

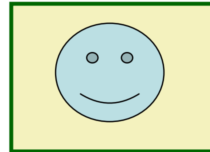

FIG. 3.4. *Original Algorithm (4/8)*

mob(s1)= t+1
loc(s1)= s2
pres(s1)=false
ack(s1)=negative

mob(s2)=t+1
loc(s2)=s2
pres(s2)=false
ack(s2)=negative

mob(s3)=t+2
loc(s3)=s3
pres(s3)=true
ack(s3)=negative

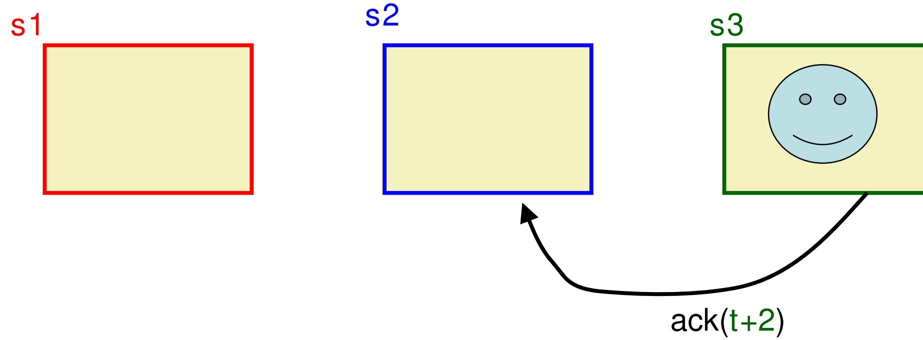

FIG. 3.5. *Original Algorithm (5/8)*

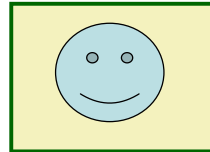

FIG. 3.6. *Original Algorithm (6/8)*

mob(s1)= t+1
loc(s1)= s2
pres(s1)=false
ack(s1)=negative

mob(s2)=t+2
loc(s2)= s3
pres(s2)=false
ack(s2)=negative

mob(s3)=t+2
loc(s3)=s3
pres(s3)=true
ack(s3)=negative

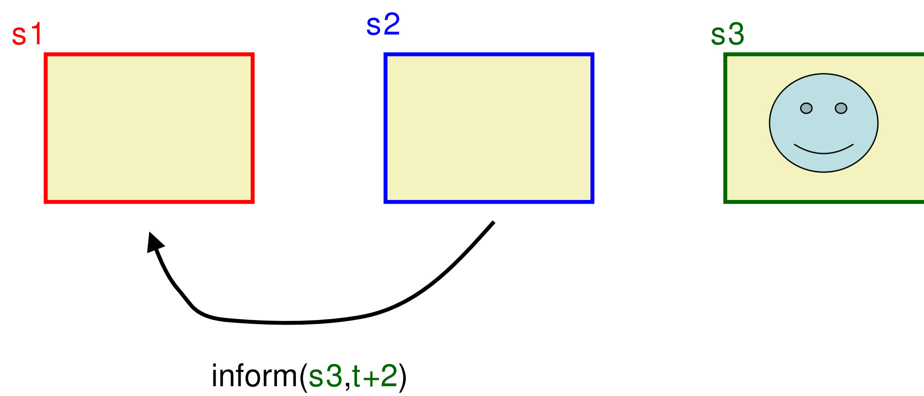

FIG. 3.7. *Original Algorithm (7/8)*

mob(s1)= t+2
loc(s1)= s3
pres(s1)=false
ack(s1)=negative

mob(s2)=t+2
loc(s2)= s3
pres(s2)=false
ack(s2)=negative

mob(s3)=t+2
loc(s3)=s3
pres(s3)=true
ack(s3)=negative

s1

s2

s3

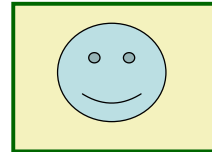

FIG. 3.8. *Original Algorithm (8/8)*

4. Fault-Tolerant Algorithm. The intuition of our solution to the problem of failures is to introduce some *redundancy* in the information that sites maintain about agents' positions. To this end, we bring three novel aspects to the original algorithm. First, agents remember N previous different sites that they have visited. Second, once an agent arrives at a new location, the agent's N previous locations are informed of its new position. Third, sites remember up to N different positions for an agent and their associated timestamps. With these changes in place, we shall establish that the algorithm is able to determine the agent's position correctly, provided that the number of stopping failures remains smaller or equal to $N-1$.

As in the original algorithm, each mobile agent is associated with a timestamp that is increased every time the agent migrates. When an agent has decided to migrate to a new location, it requests the communication layer to transport it to its new destination. When the agent arrives at a new location, an acknowledgement message containing both its new position and its newly-incremented timestamp is sent to N previous location.

We illustrate the modified algorithm by the animation in Figures 4.1 to 4.9. Each site contains a local state composed of four components, which we intuitively explain as follows:

1. *mob*: the lastest known mobility counter, i. e. timestamp, of the agent;
2. *loc*: up to N different most recently known locations visited by the agent after the current site;
3. *pres*: when the agent is present, it indicates up to N most recently visited different locations before the current site;
4. *ack*: indicating whether acknowledgement messages have to be sent or not.

A complete formalisation will follow in Section 5. In Figure 4.1, at site s_3, the agent's timestamp is $t+2$, the agent is local (and hence there is no other site visited after s_3), and the agent was known to be be in s_1 and s_2, previously at timestamps t and $t+1$, respectively.

The mobile agent located at site s_3 decides to migrate to a new location (Figure 4.2). As the agent makes the request to the transport layer to be transported and leaves s_3 (Figure 4.3), the agent state is packaged up with its previous locations s_2, s_1 and associated timestamps it had at these. The latest known timestamp and location remain unchanged at s_3 since the agent's presence elsewhere has not been confirmed yet. However, the agent is now known to be absent from s_3.

As the agent arrives at s_4 (Figure 4.5), s_4 acquires the knowledge that the agent is present locally (hence, reflected in its states *mob*, *loc* and *pres*). Furthermore, the state *ack* is changed to indicate that the new agent's position has to be communicated to its previous locations s_1 to s_3.

The sending of acknowledgement messages by s_4 (Figure 4.6) updates the *ack* state at s_4 and results in the latest known locations of the agent to be communicated to s_1, s_2, s_3 (as reflected by the change of *loc* at these sites in Figures 4.7 to 4.9).

As in the original algorithm, we allow arbitrary sites to communicate their knowledge about the agent's position and timestamp. Instead of introducing a specific message for this purpose, we overload the meaning of the acknolwedgement message (see Figures 4.7 to 4.9). This allows chains of forwarding pointers to be shortened, and hence reduce exposure to failures.

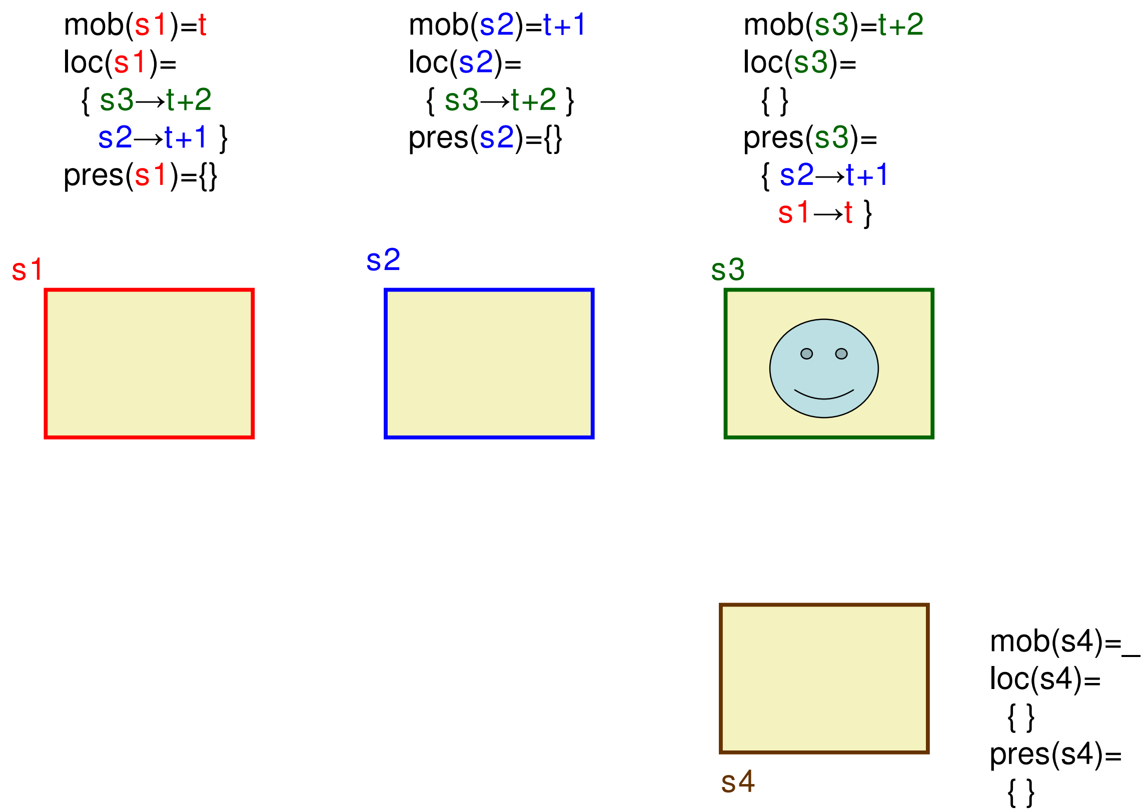

FIG. 4.1. *Agent Migration with 3-Redundancy (1/9)*

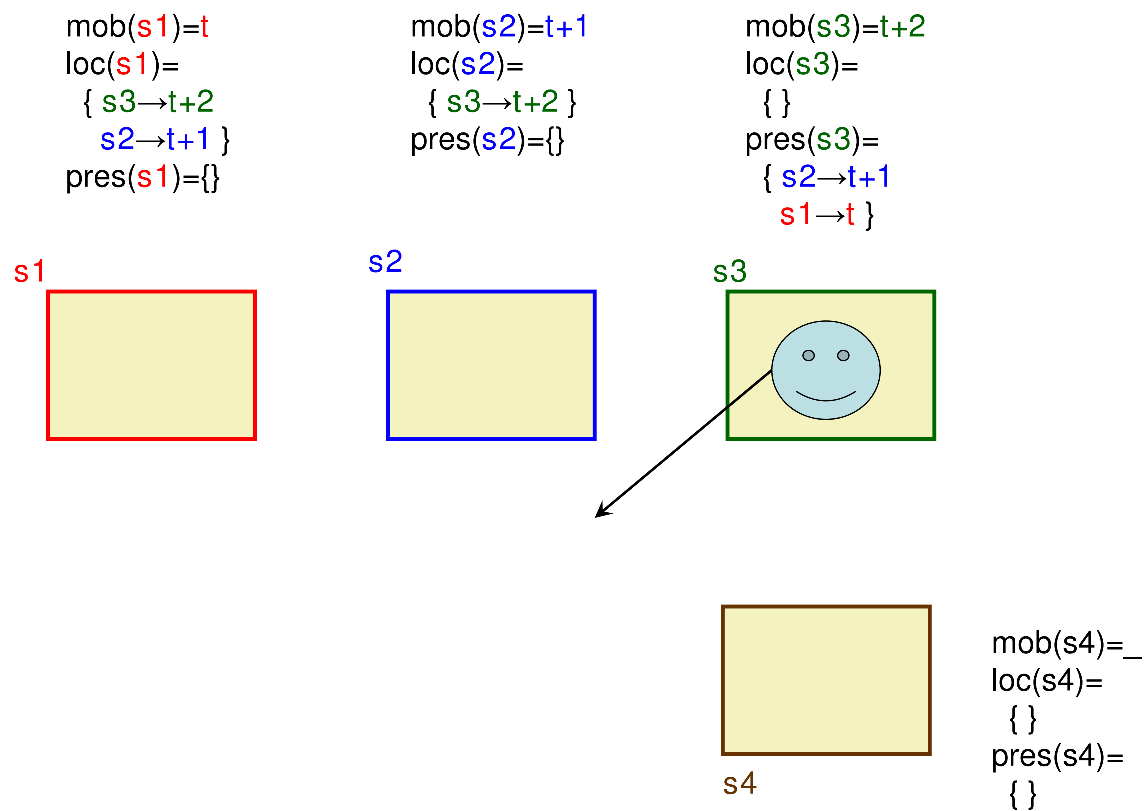

FIG. 4.2. *Agent Migration with 3-Redundancy (2/9)*

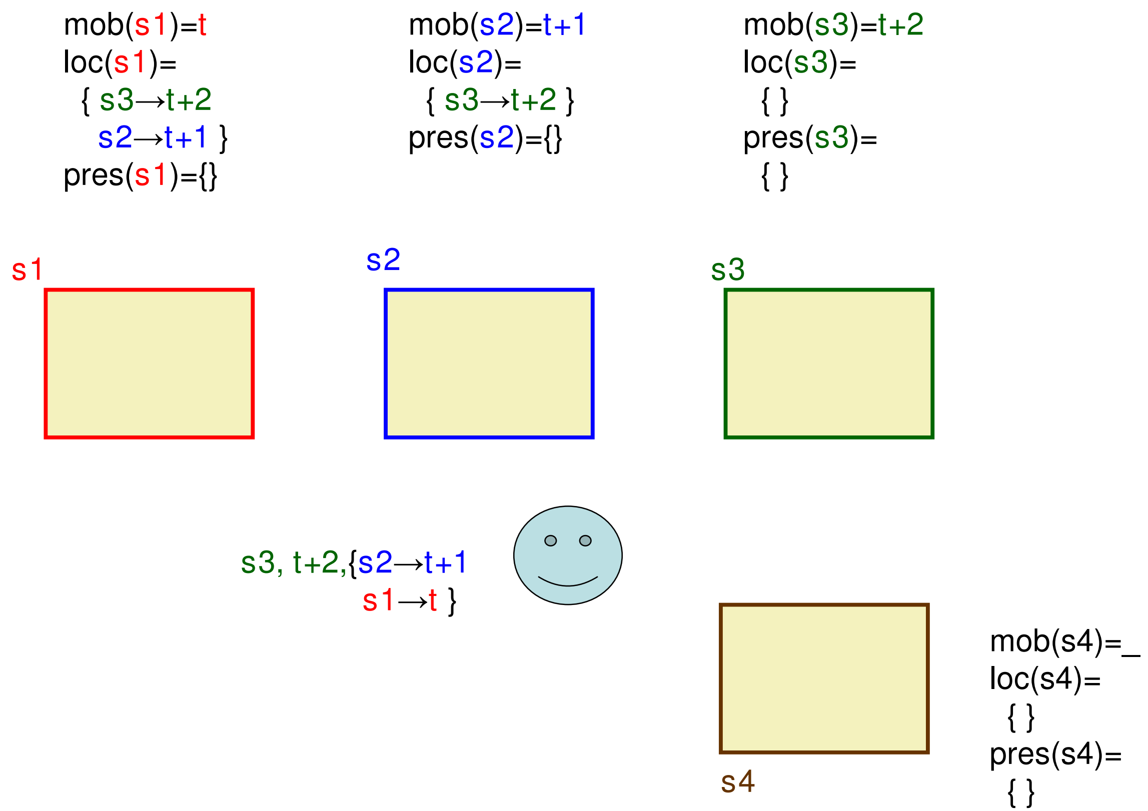

FIG. 4.3. *Agent Migration with 3-Redundancy (3/9)*

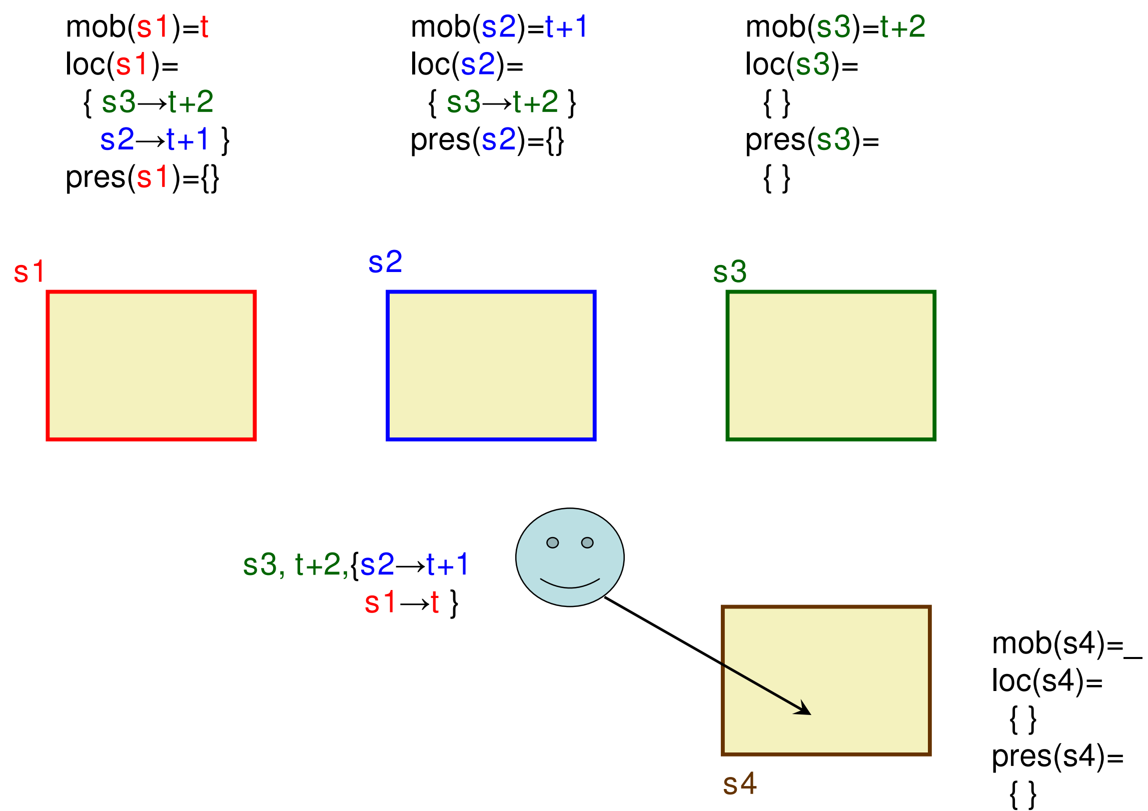

FIG. 4.4. *Agent Migration with 3-Redundancy (4/9)*

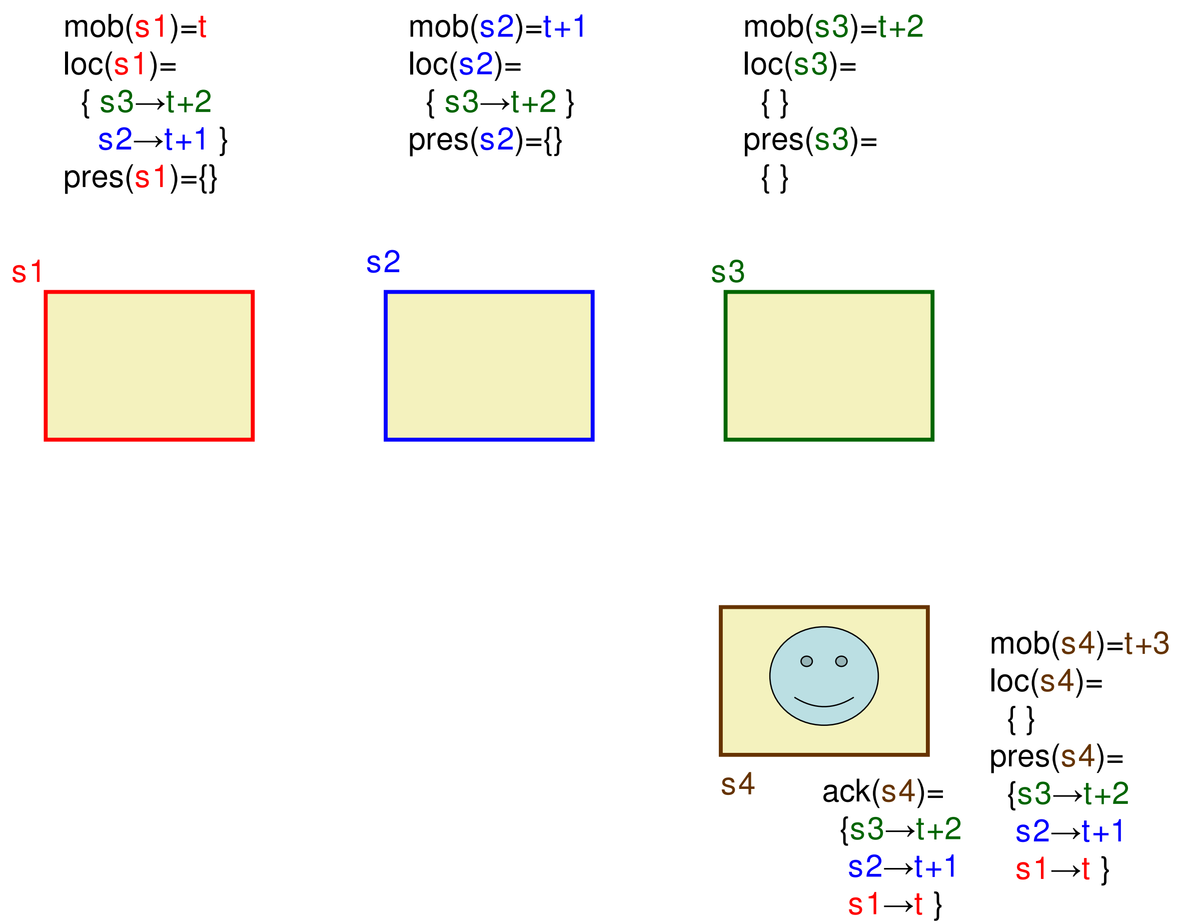

FIG. 4.5. *Agent Migration with 3-Redundancy (5/9)*

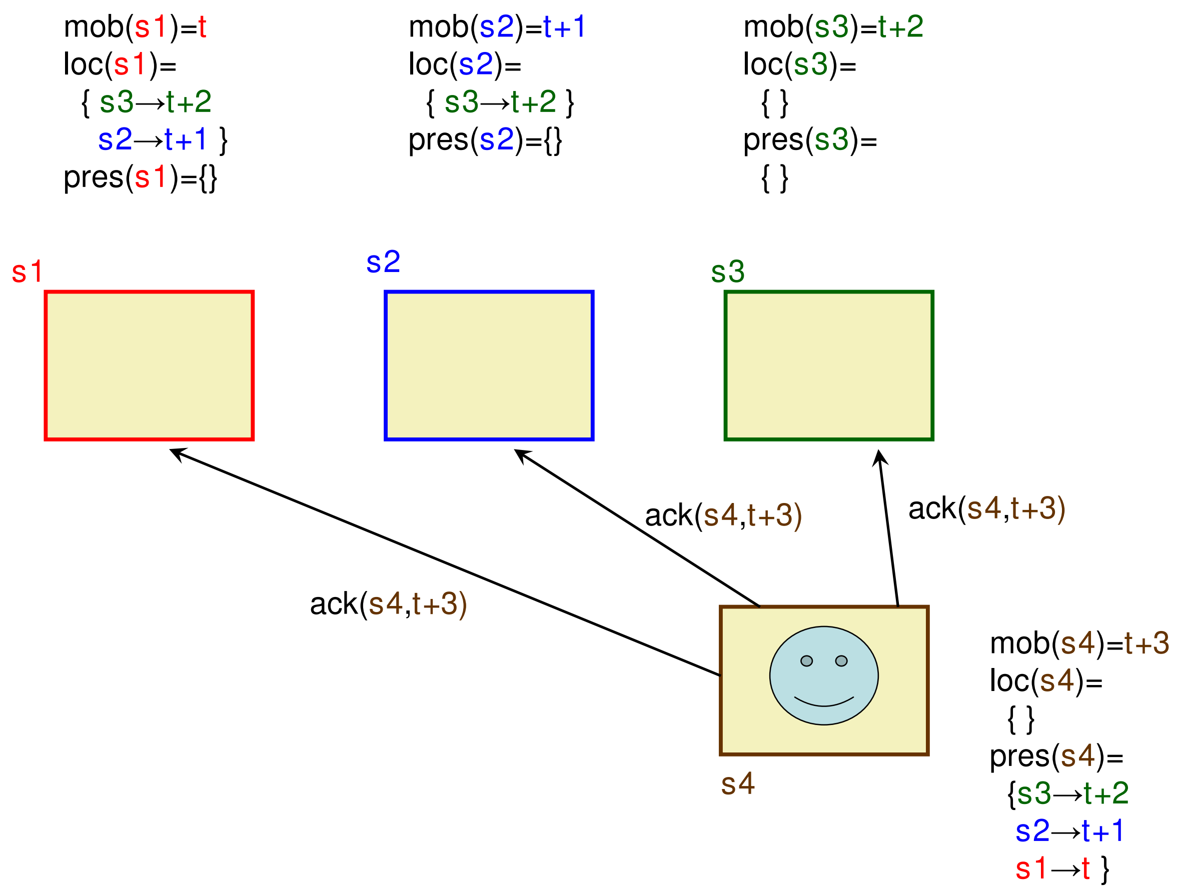

FIG. 4.6. *Agent Migration with 3-Redundancy (6/9)*

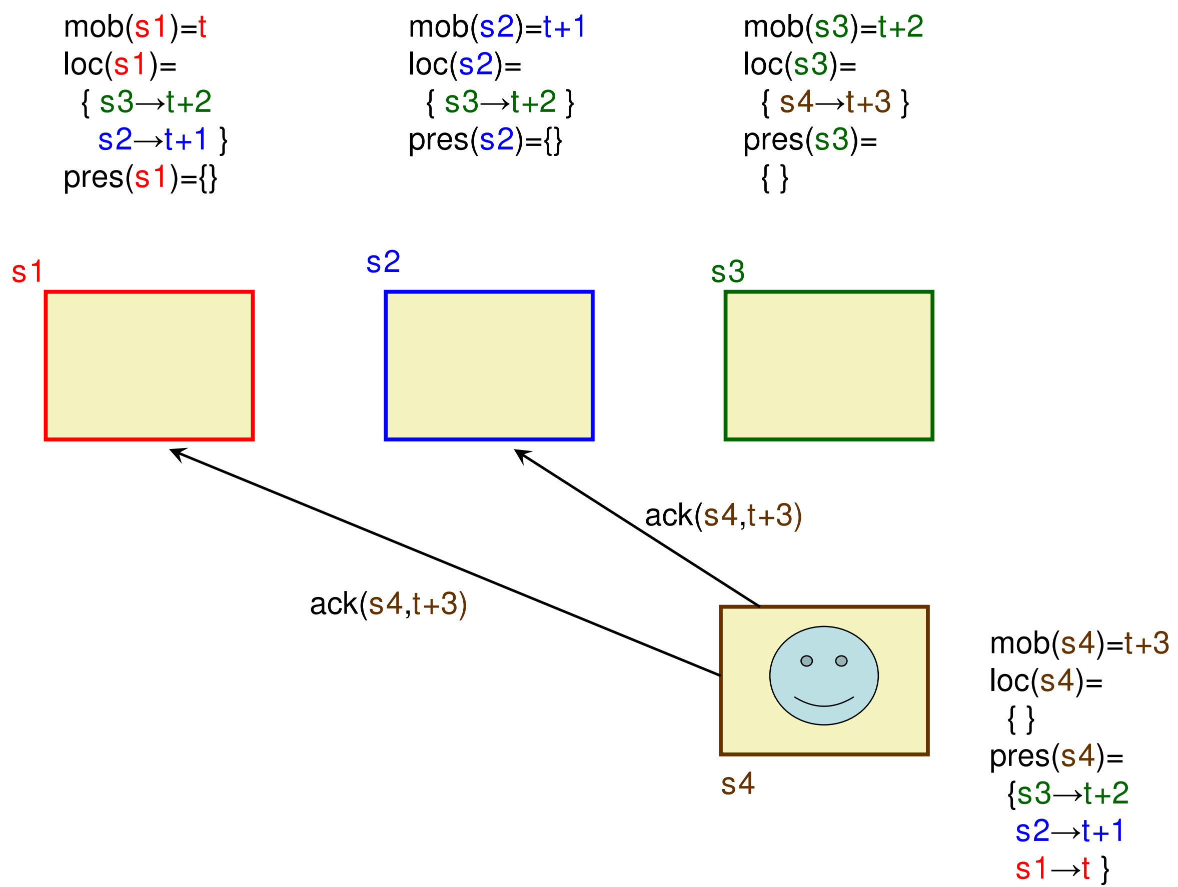

FIG. 4.7. *Agent Migration with 3-Redundancy (7/9)*

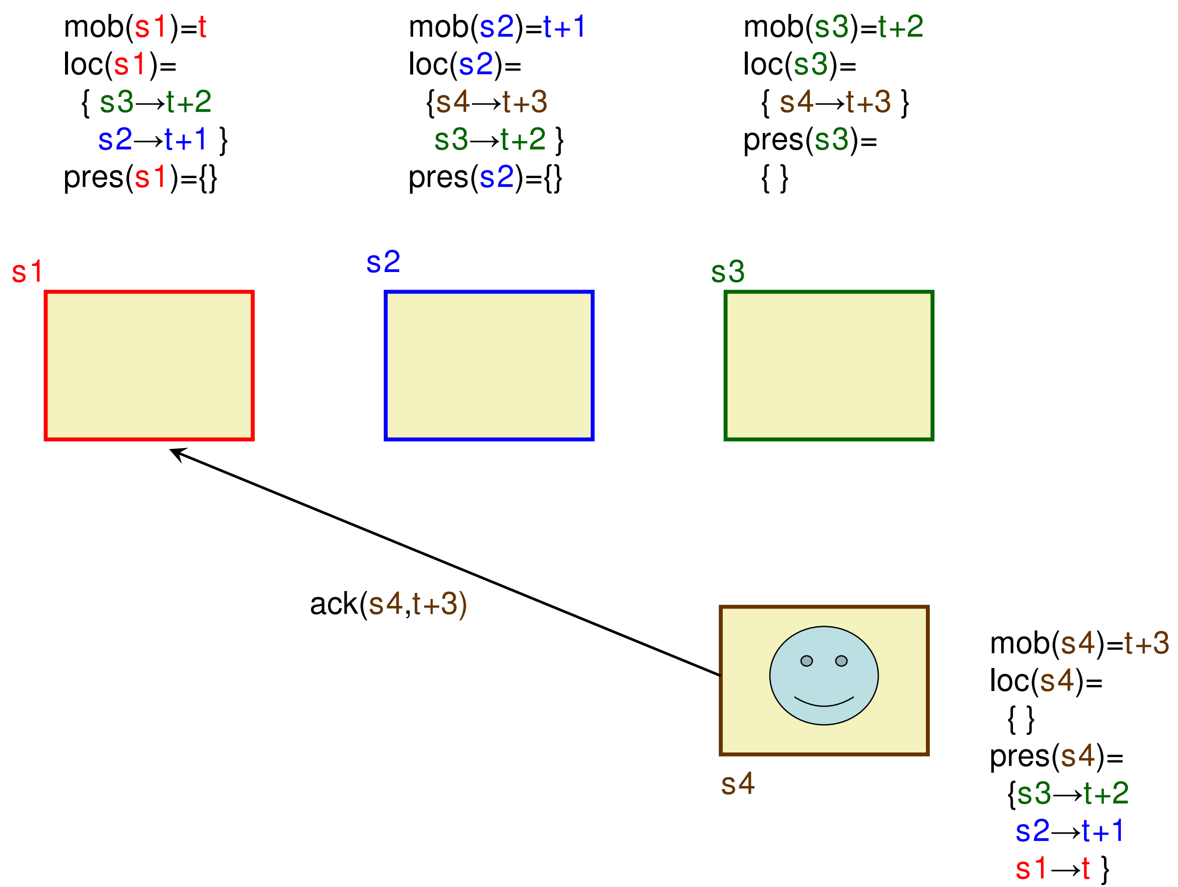

FIG. 4.8. *Agent Migration with 3-Redundancy (8/9)*

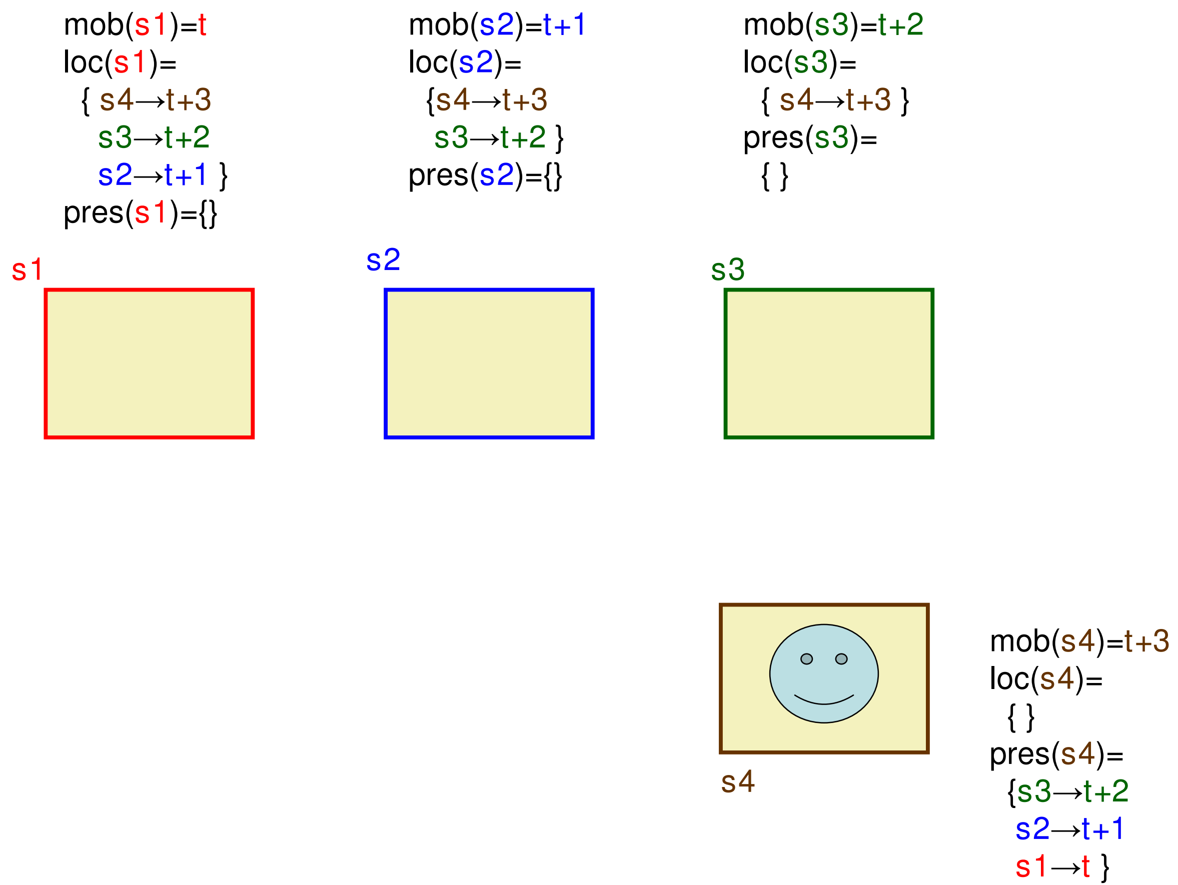

FIG. 4.9. *Agent Migration with 3-Redundancy (9/9)*

Consequently, the processing of each acknowledgement message (Figures 4.6 to 4.9), resulted in a *forwarding pointer* being set up by each recipient of these messages to the agent's latest position. Given that N acknowledgement messages are sent, N such forwarding pointers are being established, hereby introducing a redundancy in the number of pointers to the agent.

The benefit of such a redundancy is itself better illustrated by an animation, with Figures 4.10 to 4.15. Figure 4.10 is simply a replica of Figure 4.9, whereas Figure 4.11 only displays information about the latest known positions of the agent.

Such latest known agent positions are in fact forwarding pointers to locations where the agent is believed to be. Hence, in Figure 4.12, we represent such forwarding pointers graphically. We can see that from site s_1, there are many different routes that lead to s_4.

The following figures consider the presence of failures in the network. In Figure 4.13, s_2 stopped due to a failure, whereas in Figure 4.14, it is s_3 that stopped due to a failure. Among all the possible routes identified in Figure 4.12, we see in Figures 4.13 and 4.14 that there exists routes that do not involve failed nodes.

Figure 4.15 even considers two failures at s_2 and s_3, and the remaining route between s_1 and s_4 that does not involve failed nodes. Figures 4.1 to 4.9 illustrated 3-redundancy since up to three different agent locations are remembered by sites. Figures 4.15 shows that even in the present of $3-1$ failures, there is still a route to find the agent.

Remark Our purpose is to design an algorithm that is resilient to failures of *intermediary* nodes, and therefore we do not consider failures that may occur at the sender node and at the recipient node, i. e. "the first site that sends a message" and "the site where the agent is located". In other words, we are not concerned with reliability of agents themselves. Instead, systems replicating agents and using failure detectors such as [8] may be used for that purpose; they are complementary to our approach. Furthemore, in this paper, we also assume that communications are reliable.

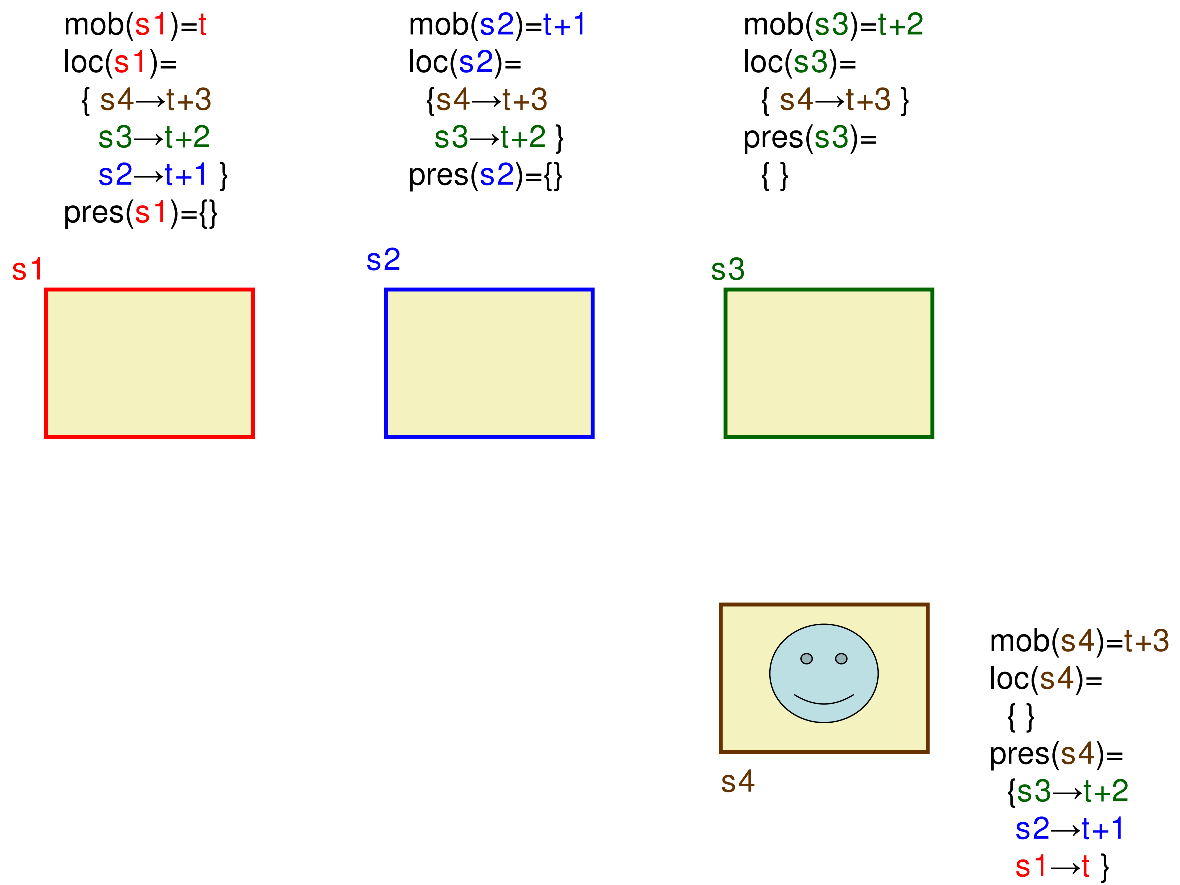

FIG. 4.10. *Alternate Paths in the Presence of Failures (1/6)*

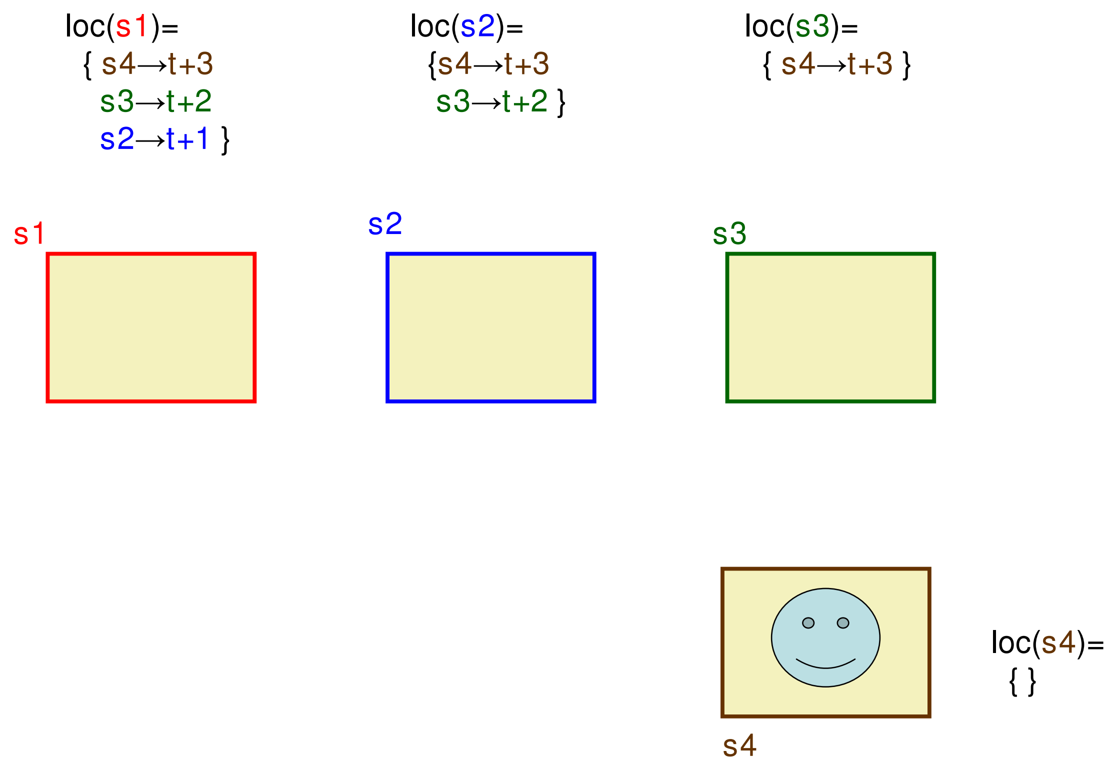

FIG. 4.11. *Alternate Paths in the Presence of Failures (2/6)*

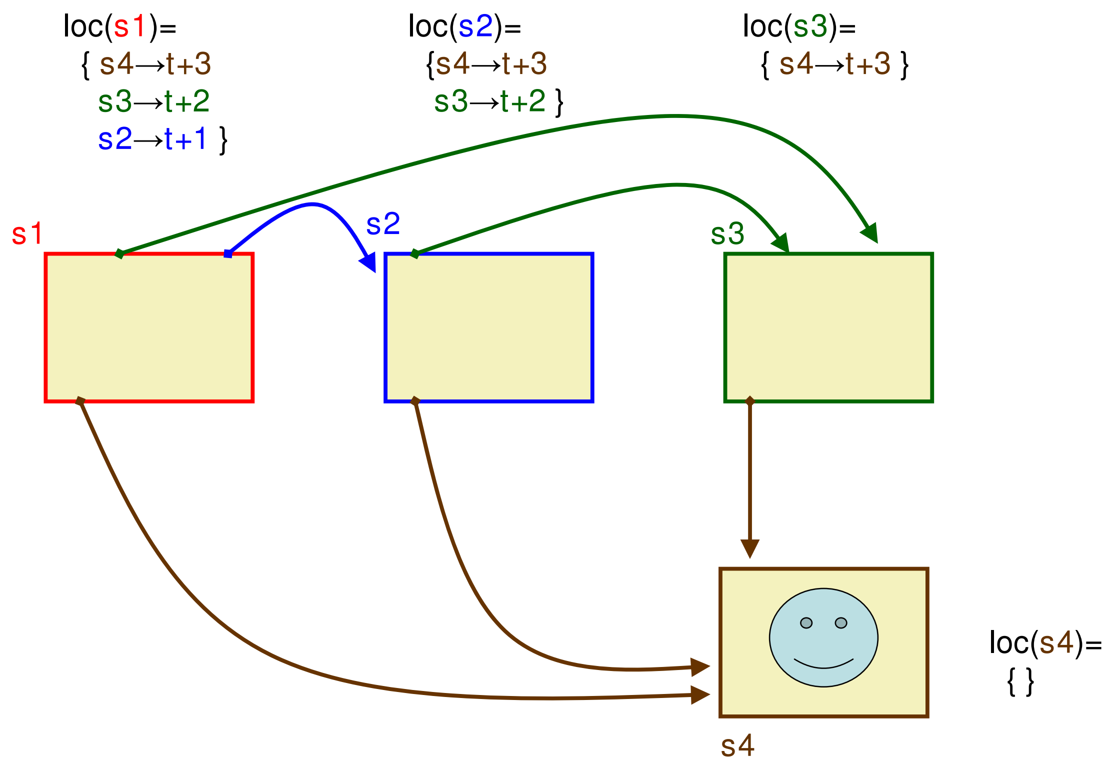

FIG. 4.12. *Alternate Paths in the Presence of Failures (3/6)*

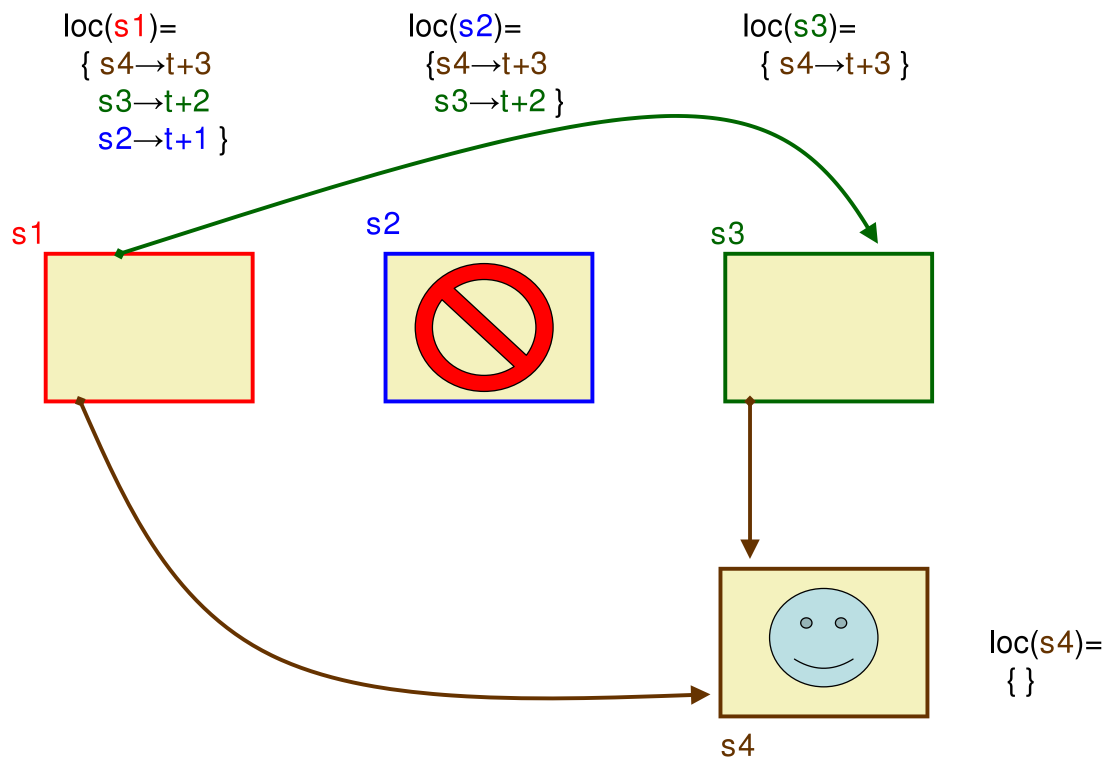

FIG. 4.13. *Alternate Paths in the Presence of Failures (4/6)*

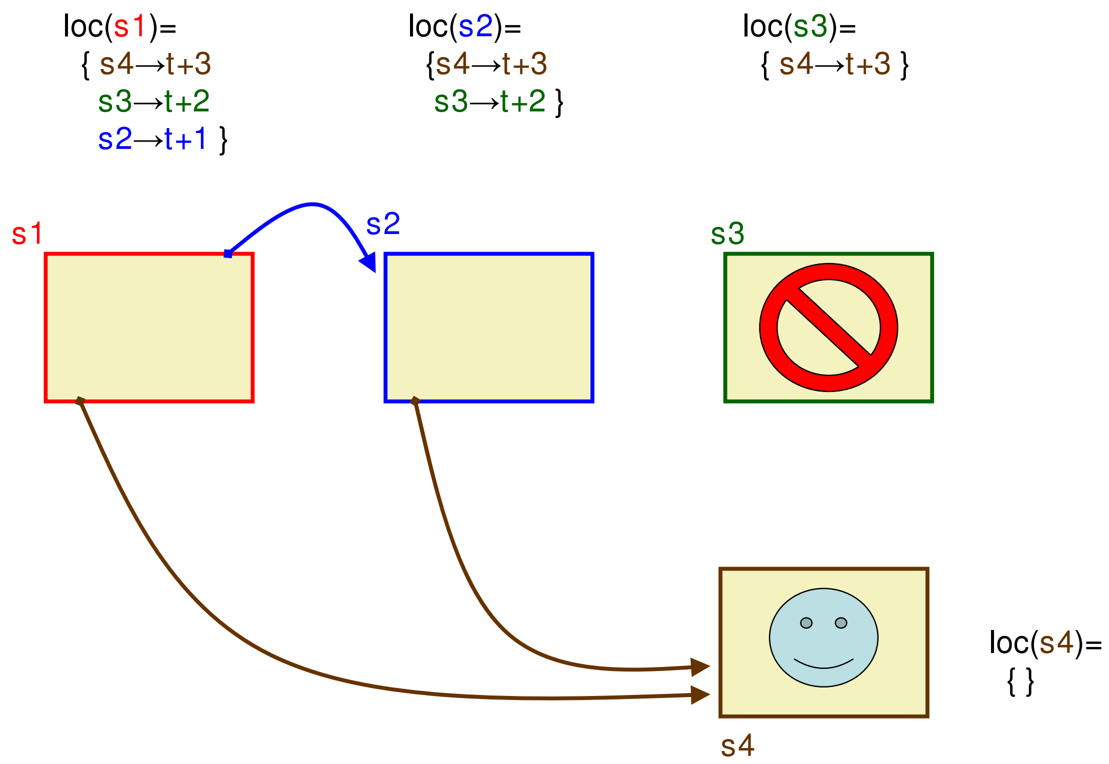

FIG. 4.14. *Alternate Paths in the Presence of Failures (5/6)*

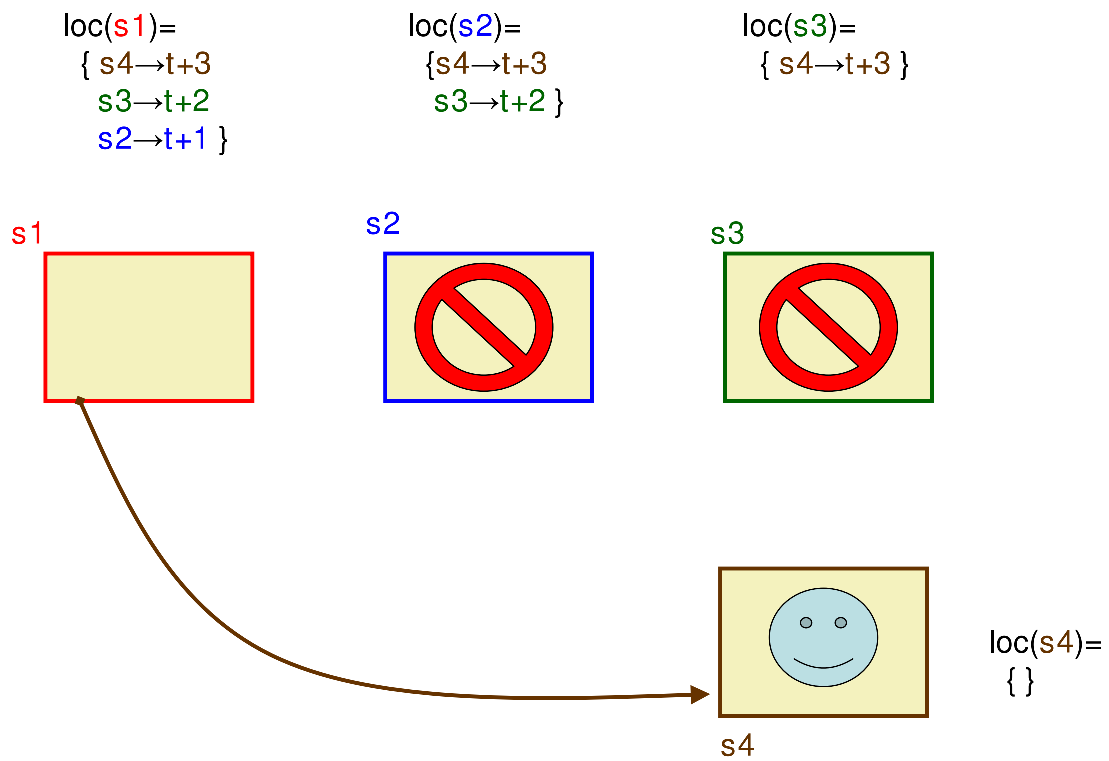

FIG. 4.15. *Alternate Paths in the Presence of Failures (6/6)*

5. Formalisation. We adopt a tried and tested framework to formalise this algorithm, which we have previously applied to several types of distributed algorithms, for mobile agents [10, 11] and for distributed reference counting [13, 12]. This formal framework has lent itself naturally to mechanical proofs in three cases.

Specifically, we model the distributed directory service as an abstract machine, whose state space is summarised in Figure 5.1. For the sake of clarity, we consider a single mobile agent; the formalisation can easily be extended to multiple agents by introducing names by which agents are being referred to. An abstract machine is composed of a set of sites taking part in a computation. Agent timestamps, which we call *mobility counters*, are defined as natural numbers. A *memory* is defined as an association list, associating locations with mobility counters; we represent an empty memory by $\emptyset$. The value N is a parameter of the algorithm. We will show that the agent's memory has a maximum size N and that the algorithm tolerates at most $N-1$ failures.

5.1. Algorithm in the Absence of Failures. The set of messages is inductively defined by two constructors. These constructors are used to construct messages, which respectively represent an agent in transit and an arrival acknowledgement. The message representing an agent in transit, typically of the form $\mathsf{agent}(s,l,\vec{M})$, contains the site s that the agent is leaving, the value l of the mobility counter it had on that site, and the agent's memory $\vec{M}$, i. e. the N previous sites it visited and associated mobility counters. The message representing an arrival acknowledgement, $\mathsf{ack}(s,l)$, contains the site s (and associated mobility counter l) where the agent is.

We assume that the network is fully connected, that communications are reliable, and that the order of messages in transit between pairs of sites is preserved. These communication hypotheses are formalised in the abstract machine by point-to-point communication links, which we define as queues using the following notations: the expression $q_1 \S q_2$ denotes the concatenation of two queues q_1, q_2, whereas $first(q)$ refers to the head of a queue q.

$$
\begin{array}{rcll}
\mathscr{S} & = & \{s_0, s_1, \ldots, s_{n_s}\} & \text{(Set of Sites)} \\
\mathscr{L} & = & \mathbb{N} & \text{(Mobility Counters)} \\
\Psi & = & list(\mathscr{S} \times \mathscr{L}) & \text{(Memory)} \\
N & \in & \mathbb{N} & \text{(Algorithm Parameter)} \\
\mathscr{M} & : & \mathsf{agent} : \mathscr{S} \times \mathscr{L} \times \Psi \rightarrow \mathscr{M} \ \mid \ \mathsf{ack} : \mathscr{S} \times \mathscr{L} \rightarrow \mathscr{M} & \text{(Messages)} \\
\mathscr{K} & = & \mathscr{S} \times \mathscr{S} \rightarrow Queue(\mathscr{M}) & \text{(Message Queues)} \\
\mathscr{LT} & = & \mathscr{S} \rightarrow \Psi & \text{(Location Tables)} \\
\mathscr{PT} & = & \mathscr{S} \rightarrow \Psi & \text{(Present Tables)} \\
\mathscr{MT} & = & \mathscr{S} \rightarrow \mathscr{L} & \text{(Mobility Counter Tables)} \\
\mathscr{AT} & = & \mathscr{S} \rightarrow \Psi & \text{(Acknowledgement Tables)} \\
\mathscr{FT} & = & \mathscr{S} \rightarrow Bool & \text{(Failure State)} \\
\mathscr{C} & = & \mathscr{LT} \times \mathscr{PT} \times \mathscr{MT} \times \mathscr{AT} \times \mathscr{FT} \times \mathscr{K} & \text{(Configurations)}
\end{array}
$$

$$
\begin{array}{ll}
\multicolumn{2}{l}{Characteristic\ variables:} \\
s \in \mathscr{S} & present_T \in \mathscr{PT} \\
m \in \mathscr{M} & mob_T \in \mathscr{MT} \\
k \in \mathscr{K} & ack_T \in \mathscr{AT} \\
c \in \mathscr{C} & fail_T \in \mathscr{FT} \\
\vec{M} \in \Psi & q \in Queue(\mathscr{M}) \\
loc_T \in \mathscr{LT} &
\end{array}
$$

FIG. 5.1. *State Space*

Each site maintains some information, which we abstract as "tables" in the abstract machine. The *location table* maps each site to a memory; for a site s, the location table indicates the sites where s believes the agent has migrated to (with their associated mobility counter). The *present table* is meant to be empty for all sites, except for the site where the agent is currently located, when the agent is not in transit; there, the present table contains the sites previously visited by the agent. The *mobility counter table* associates each site with the mobility counter the agent had when it last visited the site; the value is zero if the agent has never visited the site.

After the agent has reached a new destination, acknowledgement messages have to be sent to the N previous sites it visited. We decouple the agent's arrival from acknowledgement sending, so that transitions that deal with incoming messages are different from those that generate new messages. Consequently, we introduce a further table, the *acknowledgement table*, indicating which acknowledgements still have to be sent.

In our formalisation, we use a variable to indicate whether a machine is up and running. A site's *failure state* is allowed to change from false to true, which indicates that the site is exhibiting a failure. We are modelling *stopping failures* [7] since no transition allows a failure state to change from true to false.

A complete configuration of the abstract machine is defined as the Cartesian product of all tables and message queues. Our formalisation can be regarded as an *asynchronous* distributed system [7]. In a real implementation, tables are not shared resources, but their contents can be distributed at each site.

The behaviour of the algorithm is represented by transitions, which specify how the state of the abstract machine evolves. Figure 5.2 contain all the transitions of the distributed directory service.

For convenience, we use some notations such as *post*, *receive* or table updates, which give an imperative look to the algorithm; their definitions is as follows.

- Given a configuration $\langle loc_T, present_T, mob_T, ack_T, fail_T, k\rangle$, $mob_T(s) := V$ denotes $\langle loc_T, present_T, mob_T', ack_T, fail_T, k\rangle$, such that $mob_T'(s) = V$ and $mob_T'(s') = mob_T(s')$, $\forall\, s' \neq s$.
- A similar notation is used for other tables.
- Given a configuration, $post(s_1, s_2, m)$ denotes $\langle loc_T, present_T, mob_T, ack_T, fail_T, k'\rangle$, with $k'(s_1, s_2) = k(s_1, s_2)\S\{m\}$, and $k'(s_i, s_j) = k(s_i, s_j)$, $\forall (s_i, s_j) \neq (s_1, s_2)$.
- A similar notation is used for *receive*.

In each rule of Figure 5.2, the conditions that appear to the left-hand side of an arrow are guards that must be satisfied in order to be able to fire a transition. For instance, the first four rules contain a proposition of the form $\neg fail_T(s)$, which indicates that the rule has to occur for a site s that is up and running. The right-hand side of a rule denotes the configuration that is reached after transition. We assume that guard evaluation and new configuration construction are performed atomically.

The animation we presented in the previous section directly illustrates the applications of these rules. Figure 4.1 illustrates a configuration in which an agent has previously successively visited sites s_1, s_2, s_3 with respective timestamps $t, t+1, t+2$. In this example, we assume that the value of N is 3.

The first transition of Figure 5.2 models the actions to be performed, when an agent decides to migrate from s_1 to s_2. In the guard, we see that the present table at s_1 must be non-empty, which indicates that the agent is present at s_1. After transition, the present table at s_1 is cleared, and an agent message is posted between s_1 and s_2; the message contains the agent's origin s_1, its mobility counter $mob_T(s_1)$, and the previous content of the present table at s_1. Note that s_2, the destination of the agent, is only used to specify which communication channel the agent message must be enqueued into. The site s_1 does not need to be communicated this information, nor does it have to remember that site. In a real implementation, the agent message would also contain the complete agent code and state to be restarted by the receiver. Figure 4.3 illustrates changes in the system, after the agent has initiated its migration.

The second transition of Figure 5.2 is concerned with s_2 handling a message[1] $\mathsf{agent}(s_3, l, \vec{M})$ coming from s_1. At site s_2, tables are updated to reflect that s_2 is becoming the new agent's location, with $l+1$ its new mobility counter. Our algorithm prescribes the agent to remember N *different* sites it has visited. As s_2 may have been visited recently, we remove s_2 from $\vec{M}$, before adding the site s_3 where it was located before migration. The call $add(N, s, l, \vec{M})$ adds an association (s, l) to the memory $\vec{M}$, keeping at most N different entries with the highest timestamps. (Appendix A contains the complete definition of *add*.) In addition, the acknowledgement table of s_2 is updated, since acknowledgements have to be sent back to those previously visited sites. At this point, a proper implementation would reinstate the agent state and resume its execution. Figure 4.3 illustrates the system as an agent arrives at a new location.

According to the third transition of Figure 5.2, if the acknowledgement table on s_1 contains a pair (s_2, l_2), then an acknowledgement message $\mathsf{ack}(s_1, (mob_T(s_1)))$ has to be sent from s_1 to s_2; the acknowledgement message indicates that the agent is on s_1 with a mobility counter $mob_T(s_1)$.

According to transition receive_ack, if a site s_2 receives an acknowledgement message about site s_3 and mobility counter l, its location table has to be updated accordingly. Let us note two properties of this rule. First, we do not require the emitter s_1 of the acknowledgement message to be equal to s_3; this property allows us to overload the meaning of this message and use it for sharing information about the agent's location. Second, we make sure that updating the location

[1] Note that s_3 is not required to be equal to s_1. Indeed, we want the algorithm to be able to support sites that forward incoming agents to other sites.

For a configuration $\langle loc_T, present_T, mob_T, ack_T, fail_T, k\rangle$, legal transitions are:

$$\begin{array}{l}
\mathsf{migrate_agent}(s_1,s_2):\\
\quad s_1 \neq s_2 \ \wedge\ loc_T(s_1)=\emptyset \ \wedge\ present_T(s_1)\neq\emptyset\\
\quad \wedge\ ack_T(s_1)=\emptyset \ \wedge\ \neg fail_T(s_1)\\
\quad\quad \rightarrow \{\quad let\ \vec{M} = present_T(s_1)\\
\quad\quad\quad\quad in\quad present_T(s_1) := \emptyset\\
\quad\quad\quad\quad\quad post(s_1,s_2,\mathsf{agent}(s_1,mob_T(s_1),\vec{M}))\ \ \}
\end{array}$$

$$\begin{array}{l}
\mathsf{receive_agent}(s_1,s_2,s_3,l,\vec{M}):\\
\quad first(k(s_1,s_2)) = \mathsf{agent}(s_3,l,\vec{M})\ \wedge\ \neg fail_T(s_2)\\
\quad\quad \rightarrow \{\quad receive(s_1,s_2)\\
\quad\quad\quad\quad let\quad S' = add(N,s_3,l,remove(s_2,\vec{M}))\\
\quad\quad\quad\quad in\quad loc_T(s_2) := \emptyset\\
\quad\quad\quad\quad\quad present_T(s_2) := S'\\
\quad\quad\quad\quad\quad mob_T(s_2) := l+1\\
\quad\quad\quad\quad\quad ack_T(s_2) := S'\ \ \}
\end{array}$$

$$\begin{array}{l}
\mathsf{send_ack}(s_1,s_2,\vec{M},l_2):\\
\quad ack_T(s_1) = (s_2,l_2)\ \S\ \vec{M}\ \wedge\ \neg fail_T(s_1)\\
\quad\quad \rightarrow \{\quad ack_T(s_1) := \vec{M}\\
\quad\quad\quad\quad post(s_1,s_2,\mathsf{ack}(s_1,mob_T(s_1)))\ \ \}
\end{array}$$

$$\begin{array}{l}
\mathsf{receive_ack}(s_1,s_2,s_3,l):\\
\quad first(k(s_1,s_2)) = \mathsf{ack}(s_3,l)\ \wedge\ \neg fail_T(s_2)\\
\quad\quad \rightarrow \{\quad receive(s_1,s_2)\\
\quad\quad\quad\quad loc_T(s_2) := add(N,s_3,l,loc_T(s_2))\quad \}
\end{array}$$

$$\begin{array}{l}
\mathsf{inform}(s_1,s_2,s_3,l):\\
\quad (s_3,l) \in loc_T(s_1)\ \wedge\ \neg fail_T(s_1)\\
\quad\quad \rightarrow \{\quad post(s_1,s_2,\mathsf{ack}(s_3,l))\ \ \}
\end{array}$$

$$\begin{array}{l}
\mathsf{stop_failure}(s):\\
\quad fail_T(s) = false\\
\quad\quad \rightarrow \{\quad fail_T(s) = true\ \}
\end{array}$$

$$\begin{array}{l}
\mathsf{msg_failure}(s_1,s_2,m):\\
\quad first(k(s_1,s_2)) = m\ \wedge\ fail_T(s_2)\\
\quad\quad \rightarrow \{\quad receive(s_1,s_2)\ \}
\end{array}$$

FIG. 5.2. *Fault-Tolerant Directory Service*

table *(i)* maintains information about *different* locations, *(ii)* does not overwrite existing location information with older one. This functionality is implemented by the function *add*, whose specification required careful design to ensure correctness of the algorithm and can be found in appendix A.

According to rule inform of Figure 5.2, any site s_1 believing that the agent is located at site s_3, with a mobility counter l, may elect to communicate its belief to another site s_2. Such a belief is also communicated by an ack message. (It is here that we overload the meaning of the original ack message to allow information sharing.) It is important to distinguish the roles of the send_ack and inform transitions. The former is mandatory to ensure the correct behaviour of the algorithm, whereas the latter is optional. The purpose of inform is to propagate information about the agent's location in the system, so that the agent may be found in less steps. As opposed to previous rules, the inform rule is non-deterministic in the destination and location information in an acknowledgement message. At this level, our goal is to define a correct *specification* of an algorithm: any implementation strategy will be an instance of this specification; some of them are discussed in Section 7. Figures 4.6 illustrates the states of the system after sending acknowledgement messages, whereas Figures 4.7, 4.8, 4.9 show the effects of such messages.

5.2. Failures. The first five rules of Figure 5.2 require the site s where the transition takes place to be up and running, i. e. $\neg fail_T(s)$. Our algorithm is designed to be tolerant to *stopping failure*, according to which processes are allowed to stop somewhere in the middle of their execution [7]. We model a stopping failure by the transition stop_failure, changing the failure state of the site that exhibits the failure. Consequently, a site that has stopped will be prevented from performing any of the first five transitions of Figure 5.2.

As far as distributed system modelling is concerned, it is unrealistic to consider that messages that are in transit on a communication link remain present if the destination of the communication link exhibits a failure. Rule msg_failure shows how messages in transit to a stopped site may be lost. A similar argument may also hold for messages that were posted (but not sent yet) at a site that stops. We could add an extra rule handling such a case, but we did not do so in order to keep the number of rules limited. Thus, our communication model can be seen as using buffered inputs and unbuffered outputs.

5.3. Initial and Legal Configurations. In the initial configuration, noted c_i, we assume that the agent is at a given site *origin* with a mobility counter set to $N+1$. Obviously, at creation time, an agent cannot have visited N sites previously. Instead, the creation process elects a set $\mathscr{S}_i$ of different sites that act as "backup routers" for the agent in the initial configuration. Each site is associated with a different mobility counter in the interval $[1,N]$. Such N sites could be chosen non-deterministically by the system or could be configured manually by the user. For each site in $\mathscr{S}_i$, the location table points to the origin and to sites of $\mathscr{S}_i$ with a higher mobility counter; the location table at all other sites contains the origin and the $N-1$ first sites of $\mathscr{S}_i$. The present table at *origin* contains the sites in $\mathscr{S}_i$. A detailed formalisation of the initial configuration is available from [9]. A configuration c is said to be *legal* if there is a sequence of transitions $t_1, t_2, \ldots, t_n$ such that c is reachable from the initial configuration: $c_i \mapsto^{t_1} c_1 \mapsto^{t_2} c_2 \ldots \mapsto^{t_n} c$. We define $\mapsto^*$ as the reflexive, transitive closure of $\mapsto$.

6. Correctness. The correctness of the distributed directory service is based on two properties: safety and liveness. The *safety* of the distributed directory service ensures that it correctly tracks the mobile agent's location, in particular in the presence of failures. The *liveness* guarantees that agent location information eventually gets propagated.

We *intuitively* explain the safety property proof as follows. An acknowledgement message results in the creation of a forwarding pointer that points towards the agent's location. Forwarding pointers may be modelled by a relationship *parent* that defines a directed acyclic graph leading to the agent's location.

In the presence of failures, we show that the relationship parent contains sufficient redundancy in order to guarantee the existence of a path leading to the agent, without involving any failed site: *(i)* Sites that belong to the agent's memory have the agent's location as a parent. *(ii)* Sites that do not belong to the agent's memory have at least N parents. Consequently, if the number of failures is strictly inferior to N, each site has always at least one parent that is closer to the agent's location; by repeating this argument, we can find the agent's location.

We summarise the liveness result similar to the one in [10]. A *finite* amount of transitions can be performed from any legal configuration (if we exclude migrate_agent and inform). Furthermore, we can prove that, if there is a message at the head of a communication channel, there exists a transition of the abstract machine that consumes that message. Consequently, if we assume that message delivery and machine transitions are fair, and if the mobile agent is stationary at a location, then location tables will eventually be updated, which proves the liveness of the algorithm.

All proofs were mechanically derived using the proof assistant Coq [1]. Coq is a theorem prover whose logical foundation is constructive logic. The crucial difference between constructive logic and classical logic is that $\neg\neg p \supset p$ does not hold in constructive logic. The consequence is that the formulation of proofs and properties must make use of constructive and decidable statements. Due to space restriction, we do not include the proofs but they can be downloaded from [9]. The notation adopted in the paper is pretty-printed concise version of the mechanically established one.

7. Algorithm and Proof Discussion. The constructive proof of the initial algorithm without fault-tolerance helped us understand the different invariants that needed to be preserved. In particular, the algorithm maintains a directed acyclic graph leading to the agent's position; interestingly, short-cutting chains of pointers by propagating acknowledgement messages ensures that the graph remains connected and acyclic. Using the same mechanism of timestamp in combination with replication preserves a similar invariant in the presence of failures.

Interestingly, the fault-tolerant algorithm turned out to be simpler than its non-fault-tolerant version, because it uses less rules; furthermore, its correctness proof was easier to derive. When N is equal to 1, the algorithm has the same observable behaviour as [10]. From a practical point of view, generating the mechanical proof still remained a tedious process, though simpler, because it needed some 25000 tactic invocations, of which 5000 for the formalisation of the abstract machine were reused from our initial work.

The complexity of the algorithm is linear in N as far as the number of messages (N acknowledgement messages per migration), message length (size of a memory is $O(N)$), space per site (size of a memory is $O(N)$), and time per migration are concerned. Our proof established the correctness in the worst-case scenario. Indeed, the algorithm may tolerate more than N failures provided that one parent, at least, remains up and running for each site.

For a given application, the designer will have to choose the value of N. If N is chosen to be equal to the number of nodes in the network, the system will be fully realiable but its complexity, even though linear, is too high on an Internet scale. Instead, an engineering decision should be made: in a practical network, from network statistics, one can derive the probability of obtaining $1, 2, \ldots, N$ simultaneous failures. For each application, and for the quality of service it requires, the designer selects the appropriate failure probability, which determines the number of simultaneous failures the system should be able to tolerate.

A remarkable property of the algorithm is that it does not impose any delay upon agents when they initiate a migration. Forwarding pointers are created temporarily until a stable situation is reached and they are removed. This has to be contrasted with the home agent approach, which requires the agent to notify its homebase, before and after each migration. Interestingly, our algorithm does not preclude us also from using other algorithms; we could envision a system where algorithms are selected at runtime according to the network conditions and the quality of service requirements of the application.

Propagating agent location information with rule inform is critical in order to shorten chains of forwarding pointers, because shorter chains reduce the cost of finding an agent's location. The ideal strategy for sending these messages depends on the type of distributed system, and on the applications using the directory service. A range of solutions is possible and two extremes of the spectrum are easily identifiable. In an eager strategy, every time a mobile agent migrates, its new location is broadcasted to all other sites; such a solution is clearly not acceptable for networks such as the Internet. Alternatively, a lazy strategy could be adopted [14] but it requires cooperation with the message router. The recipient of a message may inform its emitter, when the recipient observes that that the emitter has out-of-date routing information. In such a strategy, tables are only updated when application messages are sent.

In Section 5, communication channels in the abstract machine are defined as queues. We have established that swapping any two messages in a given channel does not change the behaviour of the algorithm; in other words, messages do not need to be delivered in order.

Message Router. This paper studied a distributed directory service, and we can sketch two possible uses for message routing.

Simple Routing. The initial message router [10] can be adopted to the new distributed directory service. A site receiving a message for an agent that is not local forwards the message to the site appearing in its location table with the highest mobility counter; if the location table is empty, messages are accumulated until the table is updated. This simple algorithm does not use the redundancy provided by the directory service and is therefore not tolerant to failure.

Parallel Flooding. A site must endeavour to forward a message to N sites. If required, it has to keep copies of messages until N acknowledgements have been received. By making use of redundancy, this algorithm would guarantee the delivery of messages. We should note that the algorithm needs a mechanism to clear messages that have been delivered and are still held by intermediate nodes.

8. Further Related Work. Murphy and Picco [15] present a reliable communication mechanism for mobile agents. Their study is not concerned with nodes that exhibit failures, but with the problem of guaranteeing delivery in the presence of runaway agents. Whether their approach could be combined with ours remains an open question.

Lazar *et al.* [6] migrate mobile agents along a logical hierarchy of hosts, and also use that topology to propagate messages. As a result, they are able to give a logarithmic bound on the number of hops involved in communication.

Their mechanism does not offer any redundancy: consequently, stopping failures cannot be handled, though they allow reconnections of temporarily disconnected nodes.

Baumann and Rothermel [2] introduce the concept of a shadow as a handle on a mobile agent that allows applications to terminate a mobile agent execution by notifying the termination to its associated shadow. Shadows are also allowed to be mobile. Forwarding pointers are used to route messages to mobile agents and mobile shadows. Some fault-tolerance is provided using a mechanism similar to Jini leases, requiring message to be propagated after some timeout. This differs from our approach that relies on information replication to allow messages to be routed through multiple routes.

Mobile computing devices share with mobile agents the problem of location tracking. Prakash and Singhal [19] propose a distributed location directory management scheme that can adapt to changes in geographical distribution of mobile hosts population in the network and to changes in mobile host location query rate. Location information about mobile hosts is replicated at $O(\sqrt{m})$ base stations, where m is the total number of base stations in the system. Mobile hosts that are queried more often than others have their location information stored at a greater number of base stations. The proposed algorithm uses replication to offer improved performance during lookups and updates, but not to provide any form of fault tolerance.

9. Conclusion. In this paper, we have presented a fault-tolerant distributed directory service for mobile agents. Combined with a message router, it provides a reliable communication layer for mobile agents. The correctness of the algorithm is stated in terms of its safety and liveness.

Our formalisation is encoded in the mechanical proof assistant Coq, which we used for carrying out the proof of correctness. The constructive proof gives an ideal insight of the algorithm that can be exploited to specify a reliable message router. Overall, this would lead to a fully mechanically proven correct mobile agent system. Besides message routing, other problems deserve formalisation and mechanical proofs, such as security and authentication methods for mobile agents.

10. Acknowledgements. Thanks to Paul Groth and Zheng Chen for their feedback on the paper. This work was initiated during the QinetiQ and EPSRC sponsored Magnitude project (reference GR/N35816).

REFERENCES

[1] B. Barras, S. Boutin, C. Cornes, J. Courant, J. Filliatre, E. Giménez, H. Herbelin, G. Huet, C. M. noz, C. Murthy, C. Parent, C. Paulin, A. Saïbi, and B. Werner, *The Coq Proof Assistant Reference Manual – Version V6.1*, Tech. Report 0203, INRIA, August 1997.

[2] J. Baumann and K. Rothermel, *The shadow approach: An orphan detection protocol for mobile agents*, in Second Int. Workshop on Mobile Agents MA'98, Lecture Notes in Computer Science, Springer-Verlag, 1998, pp. 2–13.

[3] K. Bharat and L. Cardelli, *Migratory Applications*, in Mobile Object Systems: Towards the Programmable Internet, C. Tschudin and J. Vitek, eds., Springer-Verlag, Apr. 1997, pp. 131–149. Lecture Notes in Computer Science No. 1222.

[4] G. Cabri, L. Leonardi, and F. Zambonelli, *Reactive Tuple Spaces for Mobile Agent Coordination* , in Proceedings of the 2nd International Workshop on Mobile Agents (MA'98), no. 1477 in LNCS, 1998.

[5] D. B. Lange and M. Ishima, *Programming and Deploying Java Mobile Agents with Aglets*, Addison-Wesley, 1998.

[6] S. Lazar, I. Weerakoon, and D. Sidhu, *A Scalable Location Tracking and Message Delivery Scheme for Mobile Agents*, tech. report, University of Maryland, 1998.

[7] N. Lynch, *Distributed Algorithms*, Morgan Kaufmann Publishers, Dec. 1995.

[8] S. Mishra, X. Jiang, and B. Yang, *Providing Fault Tolerance to Mobile Intelligent Agents*, in Proceedings of the ISCA 8th International Conference on Intelligent Systems, Denver, June 1999.

[9] L. Moreau, *A Fault-Tolerant Distributed Directory Service for Mobile Agents: the Constructive Proof in Coq*. Available from `http://www.ecs.soton.ac.uk/~lavm/projects/algorithms/coq/failure/`, Sept. 2000.

[10] ———, *Distributed Directory Service and Message Router for Mobile Agents*, Science of Computer Programming, 39 (2001), pp. 249–272.

[11] ———, *A Fault-Tolerant Directory Service for Mobile Agents based on Forwarding Pointers*, in The 17th ACM Symposium on Applied Computing (SAC'2002) — Track on Agents, Interactions, Mobility and Systems, Madrid, Spain, Mar. 2002, pp. 93–100.

[12] L. Moreau, P. Dickman, and R. Jones, *Birrell's Distributed Reference Listing Revisited*, ACM Transactions on Programming Languages and Systems (TOPLAS), 27 (2005), p. 52.

[13] L. Moreau and J. Duprat, *A Construction of Distributed Reference Counting*, Acta Informatica, 37 (2001), pp. 563–595.

[14] L. Moreau and D. Ribbens, *Mobile Objects in Java*, Scientific Programming, 10 (2002), pp. 91–100. Special issue of the International Workshop on Performance-oriented Application Development for Distributed Architectures (PADDA'2001).

[15] A. L. Murphy and G. P. Picco, *Reliable Communication for Highly Mobile Agents*, in First International Symposium on Agent Systems and Applications/Third International Symposium on Mobile Agents (ASA/MA'99), Oct. 1999.

[16] S. Nog, S. Chawla, and D. Kotz, *An RPC mechanism for transportable agents*, Tech. Report TR96-280, Department of Computer Science, Dartmouth College, Hanover, N.H., 1996.

[17] ObjectSpace, *Voyager*. `http://www.objectspace.com/`.

[18] G. P. Picco, A. L. Murphy, and G.-C. Roman, *LIME: Linda meets mobility*, in International Conference on Software Engineering, 1999, pp. 368–377.

[19] R. PRAKASH AND M. SINGHAL, *A Dynamic Approach to Location Management in Mobile Computing Systems*, in The 8th International Conference on Software Engineering and Knowledge Engineering (SEKE'96), Lake Tahoe, Nevada, June 1996, pp. 488–495.

[20] H. K. TAN, *Interaction tracing for mobile agent security*, PhD thesis, Electronics and computer science, Southampton, UK, Mar. 2004.

[21] G. VIGNA, ed., *Mobile Agents and Security*, vol. 1419 of LNCS State-of-the-Art Survey, Springer-Verlag, June 1998.

[22] P. WOJCIECHOWSKI AND P. SEWELL, *Nomadic Pict: Language and Infrastructure Design for Mobile Agents*, in First International Symposium on Agent Systems and Applications/Third International Symposium on Mobile Agents (ASA/MA'99), Oct. 1999.

Appendix A. ADD FUNCTION. The function *add* adds a pair site–timestamp to an association list, making sure that no two entries have a same timestamp or site. A maximum of N entries is kept in the association list, and they are sorted by decreasing timestamp order.

A functional definition of *add* (close to its definition in Coq) appears below, and it uses auxiliary functions `remove` to remove an entry with a specific site from an association list and `firstN` which keeps the first N entries of an association list.

```
fun add (N:int;s1:site;n1:int;q:(Alist site int)) :=
   (firstN N (insert s1 n1 q))

fun insert (s:site;n:int;q:(Alist site int)) :=
  match q
   nil => (cons (s,l) q)
   (cons (s1,n1) q') =>
       if (s=s1)
       then
         if (n <=  n1)
         then q
         else (cons (s,n) q')
       else
         if (n<n1)
         then (cons (s1,n1) (insert s n q'))
         elif (n=n1)
         then q
         else (cons (s,n) (remove s q))
```

Edited by: Henry Hexmoor, Marcin Paprzycki, Niranjan Suri
Received: October 1, 2006
Accepted: December 12, 2006