

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCSE-2003-21

2003-04-17

LIME: A Coordination Middleware Supporting Mobility of Agents and Hosts

Amy L. Murphy, Gian Pietro Picco, and Gruia-Catalin Roman

LIME (Linda in a Mobile Environment) is a middleware supporting the development of applications that exhibit physical mobility of hosts, logical mobility of agents, or both. LIME adopts a coordination perspective inspired by work on the Linda model. The context for computation, represented in Linda by a globally accessible, persistent tuple space, is refined in LIME to transient sharing of identically-named tuple spaces carried by individual mobile units. Tuple spaces are also extended with a notion of location and programs are given the ability to react to specified states. The resulting model provides a minimalist set of abstractions that...

Read complete abstract on page 2.

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research

Recommended Citation

Murphy, Amy L.; Picco, Gian Pietro; and Roman, Gruia-Catalin, "LIME: A Coordination Middleware Supporting Mobility of Agents and Hosts" Report Number: WUCSE-2003-21 (2003). *All Computer Science and Engineering Research*.

https://openscholarship.wustl.edu/cse_research/1069

LIME: A Coordination Middleware Supporting Mobility of Agents and Hosts

Amy L. Murphy, Gian Pietro Picco, and Gruia-Catalin Roman

Complete Abstract:

LIME (Linda in a Mobile Environment) is a middleware supporting the development of applications that exhibit physical mobility of hosts, logical mobility of agents, or both. LIME adopts a coordination perspective inspired by work on the Linda model. The context for computation, represented in Linda by a globally accessible, persistent tuple space, is refined in LIME to transient sharing of identically-named tuple spaces carried by individual mobile units. Tuple spaces are also extended with a notion of location and programs are given the ability to react to specified states. The resulting model provides a minimalist set of abstractions that promise to facilitate rapid and dependable development of mobile applications. In this paper, we illustrate the model underlying LIME, provide a formal semantic characterization for the operations it makes available to the application developer, present its current design and implementation, and discuss lessons learned in developing applications that involve physical mobility.

LIME: A Coordination Middleware Supporting Mobility of Hosts and Agents

Amy L. Murphy, Gian Pietro Picco, and Gruia-Catalin Roman*

April 17, 2003

Abstract

LIME (Linda in a Mobile Environment) is a middleware supporting the development of applications that exhibit physical mobility of hosts, logical mobility of agents, or both. LIME adopts a coordination perspective inspired by work on the Linda model. The context for computation, represented in Linda by a globally accessible, persistent tuple space, is refined in LIME to transient sharing of identically-named tuple spaces carried by individual mobile units. Tuple spaces are also extended with a notion of location and programs are given the ability to react to specified states. The resulting model provides a minimalist set of abstractions that promise to facilitate rapid and dependable development of mobile applications. In this paper, we illustrate the model underlying LIME, provide a formal semantic characterization for the operations it makes available to the application developer, present its current design and implementation, and discuss lessons learned in developing applications that involve physical mobility.

1 Introduction

In the arena of modern distributed computing, mobility is emerging as a disruptive new trend that challenges fundamental assumptions across the board, from theoretical foundations to software engineering practices. Powerful social forces energized by advances in wireless communication, device miniaturization, and new software design techniques are creating a growing demand for applications that exploit and support *physical mobility* of hosts moving through space while maintaining connections with other hosts. At the same time, *logical mobility* has emerged as a novel architectural style that removes the static binding between software components and hosts, and enables run-time component migration for improved flexibility and performance. These two forms of mobility complement each other, in that logical mobility can provide the fluid software fabric necessary to cope with the high dynamicity imposed by physical mobility. Rapid development of mobile applications demands a new way of thinking and aggressive experimentation if a new set of best design practices is to emerge soon. The basic premise of this paper is that coordination technology can be extended for use in mobile computing and can offer an elegant solution to a set of difficult engineering problems.

Coordination is defined as a style of computing that emphasizes a high degree of decoupling among the computing components of an application. As initially proposed in Linda [10], this can be achieved by allowing independently developed agents to share information stored in a globally accessible, persistent, content-addressable data structure, typically implemented as a centralized tuple space. A small set of operations enabling the insertion, removal, and copying of tuples provides a simple and uniform interface to the tuple space. Temporal decoupling is achieved by dropping the requirement that the communicating parties be present at the time the communication takes place and spatial decoupling is achieved by eliminating the need for agents to be aware of each other's identity in order to communicate. A clean computational model, a high degree of decoupling, an abstract approach to communication, and a simple interface are the defining features of coordination technology. The transition to mobility requires one to revisit the basic model with a new intellectual bias. The process entails accommodating physical and logical distribution of tuples and the movement of hosts and agents through physical or logical spaces.

LIME (Linda In a Mobile Environment) is our response to the software engineering challenge posed by the advent of mobility. It defines a novel coordination-based approach to the development of mobile applications. LIME is

*A.L. Murphy (murphy@cs.rochester.edu) is with Dept. of Computer Science, University of Rochester, NY, USA. G.P. Picco (picco@elet.polimi.it) is with Dip. di Elettronica e Informazione, Politecnico di Milano, Italy. G.-C. Roman (roman@cs.wustl.edu) is with Dept. of Computer Science and Engineering, Washington University in St. Louis, MO, USA.

the first coordination model and middleware to address the need to integrate concerns having to do with physical mobility of hosts and logical mobility of agents. The LIME computational model assumes a set of hosts that act as containers in which agents are located. Physical connectivity among the hosts is supported by wired or wireless links and may be altered by mobility or by explicit connection and disconnection. Agents can move from one host to another reachable host of their own volition. LIME preserves the essence of the Linda model, its simplicity and decoupled style of computing, by continuing to channel all coordination actions through a simple interface perceived by each agent to be merely a local tuple space. Access to the tuple space is carried out using an extended set of tuple space operations that includes several novel constructs designed to facilitate flexible and timely responses to changes in the contents of the tuple space. Each agent may own multiple tuple spaces that may be shared with other agents within communication range. Sharing is made manifest by logically extending the contents of each tuple space to include the tuples present in all participating tuple spaces. The set of tuples being shared changes over time as a result of the agents' local control regarding sharing and in response to the mobility of both agents and hosts. When hosts come into communication range the set of shared tuple spaces expands and when they move apart it contracts. The net result is a transparently managed context that expresses itself in terms of changes in the contents of what otherwise appears to be local tuple spaces. The agent behavior is altered both by the availability of new data and by its reactive responses to contextual changes.

A distinctive feature of the effort leading to the development of LIME is the close interplay between formal semantic definition, implementation pragmatics, and application-driven evaluation of the resulting model and middleware. The insistence on formalizing the model and the semantics of the API is rooted in the conviction that precise semantics are key to dependable development, particularly when working in a novel and demanding setting such as mobility. Implementation considerations led to weakening of certain constructs, to the introduction of features the formal framework did not identify, and to the enhancement of the LIME middleware so as to ensure its applicability in a wide range of physical and logical mobile settings. Overall, the focus on application development and continuous empirical evaluation contributed to practically minded additions to the model. Ultimately, the effort culminated in a Java-based implementation of the LIME middleware [17], currently available [34] as an evolving, open source project. Several application development exercises with programmers possessing varying skill levels reinforced our conviction that a properly tailored model of coordination, such as LIME, can be an effective software engineering tool in the mobile setting.

The remainder of this paper summarizes our experience with LIME from the basic model to applications built on the middleware. Section 2 provides an informal overview of LIME. Section 3 introduces the formal semantic definition of the LIME operations. Section 4 explains our implementation strategy, and Section 5 reviews our experience with several applications developed using LIME. Finally, Section 6 discusses lessons learned and related work, and Section 7 ends the paper with some brief concluding remarks.

2 LIME: Linda in a Mobile Environment

The LIME model [24] aims at identifying a coordination layer that can be exploited successfully for designing applications that exhibit logical mobility, physical mobility, or both. The design criteria underlying LIME come from the realization that the problem of designing applications involving mobility can be regarded as a coordination problem [31], and that a fundamental issue to be tackled is the provision of good abstractions for dealing with, and exploiting, a dynamically changing context. To achieve its goal, LIME borrows and adapts the communication model made popular by Linda [10]. After presenting a concise Linda primer, the remainder of this section discusses how the core concepts of Linda are reshaped in the LIME model and embodied in the programming interface of the corresponding middleware implementation.

2.1 Linda in a Nutshell

In Linda, processes communicate through a shared *tuple space* that acts as a repository of elementary data structures, or *tuples*. A tuple space is a multiset of tuples that can be accessed concurrently by several processes. Each tuple is a sequence of typed fields, such as $\langle \text{"foo"}, 9, 27.5 \rangle$, and contains the information being communicated.

Tuples are added to a tuple space by performing an $\text{out}(t)$ operation, and can be removed by executing $\text{in}(p)$. Tuples are anonymous, thus their selection takes place through pattern matching on the tuple content. The argument p is often called a *template* or *pattern*, and its fields contain either *actuals* or *formals*. Actuals are values; the fields of the previous tuple are all actuals, while the last two fields of $\langle \text{"foo"}, ?\text{integer}, ?\text{float} \rangle$ are formals. Formals act like

“wild cards”, and are matched against actuals when selecting a tuple from the tuple space. For instance, the template above matches the tuple defined earlier. If multiple tuples match a template, the one returned by `in` is selected non-deterministically. Tuples can also be read from the tuple space using the non-destructive `rd(p)` operation. Both `in` and `rd` are blocking, i.e., if no matching tuple is available in the tuple space the process performing the operation is suspended until a matching tuple becomes available. A typical extension to this synchronous model is the provision of a pair of asynchronous primitives `inp` and `rdp`, called *probes*, that allow non-blocking access to the tuple space¹. Moreover, some variants of Linda (e.g., [33]) provide also *bulk operations*, which can be used to retrieve all matching tuples in one step. In LIME we provide a similar functionality through the `ing` and `rdg` operations, whose execution is asynchronous like in the case of *probes*².

2.2 The LIME Model

Linda characteristics resonate with the mobile setting. In particular, communication in Linda is decoupled in *time* and *space*, i.e., senders and receivers do not need to be available at the same time, and mutual knowledge of their identity or location is not necessary for data exchange. This form of decoupling is of paramount importance in a mobile setting, where the parties involved in communication change dynamically due to their migration or connectivity patterns. Moreover, the notion of tuple space provides a straightforward and intuitive abstraction for representing the computational context perceived by the communicating processes. On the other hand, decoupling is achieved thanks to the properties of the Linda tuple space, namely its global accessibility to all the processes, and its persistence—properties that are clearly hard if not impossible to maintain in a mobile environment.

2.2.1 The Core Idea: Transparent Context Maintenance

In Linda, the data accessible through the tuple space represents the data *context* available during process interaction. In the model underlying LIME, the shift from a fixed context to a dynamically changing one is accomplished by breaking up the Linda tuple space into many tuple spaces, each permanently associated to a mobile unit, and by introducing rules for transient sharing of these individual tuple spaces based on connectivity.

The individual tuple space permanently and exclusively attached to a mobile unit is referred to as the *interface tuple space* (ITS) because it provides the only access to the data context for that mobile unit. Each ITS contains the tuples the mobile unit is willing to make available to other units, and access to this data structure uses standard Linda operations, whose semantics remain basically unaffected. These tuples represent the only context accessible to a mobile unit when it is alone.

When multiple mobile units are able to communicate, either directly or transitively, we say these units form a LIME *group*. We can restrict the notion of group membership beyond simple communication, but for the purposes of this paper, we consider only connectivity. Conceptually, the contents of the ITSs of all group members are merged, or transiently shared, to form a single, large context which is accessed by each unit through its own ITS. The sharing itself is transparent to each mobile unit, however as the members of the group change, the content of the tuple space each member perceives through operations on the ITS changes in a transparent way.

The joining of a group by a mobile unit, and the subsequent merging of its local context with the group context is referred to as *engagement*, and is performed as a single, atomic operation. A mobile unit leaving a group triggers *disengagement*, that is, the atomic removal of the tuples representing its local context from the remaining group context. In general, whole groups can merge, and a group can split into several groups due to changes in connectivity.

In LIME, agents may have multiple ITSs distinguished by a name since this is recognized [7] as a useful abstraction to separate related application data. The sharing rule in the case of multiple tuple spaces relies on tuple space names: only identically-named tuple spaces are transiently shared among the members of a group. Thus, for instance, when an agent *a* owning a single tuple space named *X* joins a group constituted by an agent *b* that owns two tuple spaces named *X* and *Y*, only *X* becomes shared between the two agents. Tuple space *Y* remains accessible only to *b*, and potentially to other agents owning *Y* that may join the group later on.

Transient sharing of the ITS constitutes a very powerful abstraction, as it provides a mobile unit with the illusion of a local tuple space that contains all the tuples coming from all the units belonging to the group, without any need to

¹Additionally, Linda implementations often include also an `eval` operation which provides dynamic process creation and enables deferred evaluation of tuple fields. For the purposes of this work, however, we do not consider this operation further.

²Hereafter we often do not mention this pair of operations, since they are useful in practice but do not add significant complexity either to the model or to the implementation.




Figure 1: Transiently shared tuple spaces encompass physical and logical mobility.

know the members explicitly. The notion of transiently shared tuple space is a natural adaptation of the Linda tuple space to a mobile environment. When physical mobility is involved, and especially in the radical setting defined by mobile ad hoc networking, there is no stable place to store a persistent tuple space. Connections among machines come and go and the tuple space must be partitioned in some way. Analogously, in the scenario of logical mobility, maintaining locality of tuples with respect to the agent they belong to may be complicated. LIME enforces an a priori partitioning of the tuple space in subspaces that get transiently shared according to precise rules, providing a tuple space abstraction that depends on connectivity.

2.2.2 Encompassing Physical and Logical Mobility

In LIME, mobile hosts are connected when a communication link is available. Availability may depend on a variety of factors, including quality of service, security considerations, or connection cost; all of which can be represented in LIME, although in this paper we limit ourselves to availability determined by the presence of a functioning link. Mobile agents are connected when they are co-located on the same host, or they reside on hosts that are connected. Changes in connectivity among hosts depend only on changes in the physical communication links. Connectivity among mobile agents may depend also on arrival and departure of agents, with creation and termination of mobile agents being regarded as a special case of connection and disconnection, respectively. Figure 1 depicts the model adopted by LIME. Mobile agents are the only active components; mobile hosts are mainly roaming containers which provide connectivity and execution support for agents. In other words, mobile agents are the only components that carry a “concrete” tuple space with them.

The transiently shared ITSS belonging to multiple agents co-located on a host define a *host-level tuple space*. The concept of transient sharing can also be applied to the host-level tuple spaces of connected hosts, forming a *federated tuple space*. When a federated tuple space is established, a query on the ITS of an agent returns a tuple that may belong to the tuple space carried by that agent, to a tuple space belonging to a co-located agent, or to a tuple space associated with an agent residing on some remote, connected host.

In this model, physical and logical mobility are separated in two different tiers of abstraction. Nevertheless, many applications do not need both forms of mobility, and straightforward adaptations of the model are possible. For instance, applications that do not exploit mobile agents but run on a mobile host can employ one or more stationary agents, i.e., programs that do not contain migration operations. In this case, the design of the application can be modeled in terms of mobile hosts whose ITS is a fixed host-level tuple space. Applications that do not exploit physical mobility—and do not need a federated tuple space spanning different hosts—can exploit only the host-level tuple space as a local communication mechanism among co-located agents.

Nevertheless, it is interesting to note how mobility is not dealt with directly in LIME, i.e., there are no constructs for triggering the mobility of agents or hosts. Instead, the effect of migration is made indirectly manifest to the model and middleware only through the changes observed in the connectivity among components. This choice, that sets the nature of mobility aside, keeps our model as general as possible and, at the same time, enables different instantiations of the model based on different notions of connectivity.

2.2.3 Controlling Context Awareness

Thus far, LIME appears to foster a style of coordination that reduces the details of distribution and mobility to content changes in what is perceived as a local tuple space. This view is very powerful, and has the potential for greatly

Current location	Destination location	Defined projection
unspecified	unspecified	Entire federated tuple space
unspecified	λ	Tuples in the federated tuple space and destined to λ
ω	unspecified	Tuples in ω 's tuple space
Ω	unspecified	Tuples in Ω 's host-level tuple space, i.e., belonging to any agent at Ω
ω	λ	Tuples in ω 's tuple space and destined to λ
Ω	λ	Tuples in Ω 's host-level tuple space and destined to λ

Table 1: Accessing different portions of the federated tuple space by using location parameters. In the table, ω and λ are agent identifiers, while Ω is a host identifier.

simplifying application design in many scenarios by relieving the designer from the chore of maintaining explicitly a view of the context consistent with changes in the configuration of the system. On the other hand, this view may hide too much in domains where the designer needs more fine-grained control over the portion of the context that needs to be accessed. For instance, the application may require control over the agent responsible for holding a given tuple, something that cannot be specified only in terms of the global context. Also, performance and efficiency considerations may come into play, as in the case where application information would enable access aimed at a specific host-level tuple space, thus avoiding the greater overhead of a query spanning the whole federated tuple space. Such fine-grained control over the context perceived by the mobile unit is provided in LIME by extending the Linda operations with tuple location parameters that operate on user-defined projections of the transiently shared tuple space. Further, all tuples are implicitly augmented with two fields, representing the tuple's *current* and *destination location*. The current location identifies the single agent responsible for holding the tuple when all agents are disconnected, and the destination location indicates the agent with whom the tuple should eventually reside.

The $\text{out}[\lambda]$ operation extends out with a location parameter representing the identifier of the agent responsible for holding the tuple. The semantics of $\text{out}[\lambda](t)$ involve two steps. The first step is equivalent to a conventional $\text{out}(t)$, the tuple t is inserted in the ITS of the agent calling the operation, say ω . At this point the tuple t has a current location ω , and a destination location λ . If the agent λ is currently connected, the tuple t is moved to the destination location in the same atomic step. On the other hand, if λ is currently disconnected the tuple remains at the current location, the tuple space of ω . This “*misplaced*” tuple, if not withdrawn³, will remain misplaced unless λ becomes connected. In the latter case, the tuple will migrate to the tuple space associated with λ as part of the engagement. By using $\text{out}[\lambda]$, the caller can specify that the tuple is supposed to be placed within the ITS of agent λ . This way, the default policy of keeping the tuple in the caller's context until it is withdrawn can be overridden, and more elaborate schemes for transient communication can be developed.

Variants of the in and rd operations that allow location parameters are allowed as well. These operations, of the form $\text{in}[\omega, \lambda](p)$ and $\text{rd}[\omega, \lambda](p)$, enable the programmer to refer to a projection of the current context defined by the value of the location parameters, as illustrated in Table 1. The current location parameter enables the restriction of scope from the entire federated tuple space (no value specified) to the tuple space associated to a given host or even a given agent. The destination location is used to identify misplaced tuples.

2.2.4 Reacting to Changes in Context

In the fluid scenario we target, the set of available data, hosts, and agents change rapidly according to the reconfiguration induced by mobility. Reacting to changes constitutes a significant fraction of an application's activities. At first glance, the Linda model would seem sufficient to provide some degree of reactivity by representing relevant events as tuples, and by using the in operation to execute the corresponding reaction as soon as the event tuple appears in the tuple space. Nevertheless, in practice this solution has a number of drawbacks. For instance, programming becomes cumbersome, since the burden of implementing a reactive behavior is placed on the programmer rather than the system. Moreover, enabling an asynchronous reaction would require the execution of in in a separate thread of control, hence degrading performance. Therefore, LIME explicitly extends the basic Linda tuple space with the notion of *reaction*. A reaction $\mathcal{R}(s, p)$ is defined by a code fragment s that specifies the actions to be executed when a tuple matching the pattern p is found in the tuple space. The semantics of reactions are based on the Mobile UNITY

³Note how specifying a destination location λ implies neither guaranteed delivery nor ownership of the tuple t to λ . Linda rules for non-deterministic selection of tuples are still in place; thus, it might be the case that some other agent may withdraw t from the tuple space before λ , even after t reached λ 's ITS.

reactive statements [15], described formally in a later section. Informally, a reaction can *fire* if a tuple matching pattern p exists in the tuple space. After every regular tuple space operation, a reaction is selected non-deterministically and, if it is enabled, the statements in s are executed in a single, atomic step. This selection and execution continues until no reactions are enabled, at which point normal processing resumes. Blocking operations are not allowed in s , as they may prevent the execution of s from terminating.

LIME reactions can be explicitly registered and deregistered on a tuple space, and hence do not necessarily exist throughout the life of the system. Moreover, a notion of *mode* is provided to control the extent to which a reaction is allowed to execute. A reaction registered with mode ONCE is allowed to fire only one time, i.e., after its execution it becomes automatically deregistered, and hence removed from the reactive program. Instead, a reaction registered with mode ONCEPERTUPLE is allowed to fire an arbitrary number of times, but never twice for the same tuple. Finally, reactions can be annotated with location parameters, with the same meaning discussed earlier for in and rd. Hence, the full form of a LIME reaction is $\mathcal{R}[\omega, \lambda](s, p, m)$, where m is the mode.

Reactions provide the programmer with very powerful constructs. They enable the specification of the appropriate actions that need to take place in response to a *state* change and allow their execution in a single atomic step. In particular, it is worth noting how this model is much more powerful than many event-based ones [32], including those exploited by tuple space middleware such as TSpaces [12] and JavaSpaces [13], that are typically stateless and provide no guarantee about the atomicity of event reactions.

Nevertheless, this expressive power comes at a price. In particular, when multiple hosts are present, the content of the federated tuple space depends on the content of the tuple spaces belonging to physically distributed, remote agents. Thus, maintaining the requirements of atomicity and serialization imposed by reactive statements requires a distributed transaction encompassing several hosts for every tuple space operation on any ITS—very often, an impractical solution. For specific applications and scenarios, e.g., those involving a very limited number of nodes, these kind of reactions, referred to as *strong reactions*, would still be reasonable and therefore they remain part of the model. For practical performance reasons, however, our implementation currently limits the use of strong reactions by restricting the current location field to be a host or agent, and by enabling a reaction to fire only when the matching tuple appears on the same host as the agent that registered the reaction. As a consequence, a mobile agent can register a reaction for a host different from the one where it is residing, but such a reaction remains disabled until the agent migrates to the specified host. These constraints effectively force the *detection* of a tuple matching p and the corresponding *execution* of the code fragment s to take place (atomically) on a single host, and hence does not require a distributed transaction.

To strike a compromise between the expressive power of reactions and the practical implementation concerns, we introduce a new reactive construct that allows some form of reactivity spanning the whole federated tuple space but with weaker semantics. The processing of a *weak reaction* proceeds as in the case of a strong reaction, but detection and execution do not happen atomically: instead, execution is guaranteed to take place only eventually, after a matching tuple is detected. The execution of s takes place on the host of the agent that registered the reaction.

2.2.5 Exposing System Configuration

It is interesting to note that the extension of Linda operations with location parameters, as well as the other operations discussed thus far, foster a model that hides completely the details of the system (re)configuration that generated those changes. For instance, if the probe $\text{inp}[\omega, \lambda](p)$ fails, this simply means that no tuple matching p is available in the projection of the federated tuple space defined by the location parameters $[\omega, \lambda]$. It cannot be directly inferred whether the failure is due to the fact that agent ω does not have a matching tuple, or simply agent ω is currently not part of the group.

Without awareness of the system configuration, only a partial context awareness can be accomplished, where applications are aware of changes in the portion of context concerned with application data. Although this perspective is often enough for many mobile applications, in many others the portion of context more closely related to the system configuration plays a key role. For instance, a typical problem is to react to departure of a mobile unit, or to determine the set of units currently belonging to a LIME group. Interestingly, LIME provides this form of awareness of the system configuration by using the same abstractions discussed thus far: through a transiently shared tuple space conventionally named LimeSystem to which all agents are permanently bound. The tuples in this tuple space contain information about the mobile units present in the group and their relationship, e.g., which tuple spaces they are sharing or, for mobile agents, which host they reside on. Insertion and withdrawal of tuples in LimeSystem is a prerogative of the run-time support. Nevertheless, applications can read tuples and register reactions to respond to

```

public class LimeTupleSpace {
    public LimeTupleSpace(String name);
    public String getName();
    public boolean isOwner();
    public boolean isShared();
    public boolean setShared(boolean isShared);
    public static boolean setShared(LimeTupleSpace[] lts, boolean isShared);
    public void out(ITuple tuple);
    public ITuple in(ITuple template);
    public ITuple rd(ITuple template);
    public void out(AgentLocation destination, ITuple tuple);
    public ITuple in(Location current, AgentLocation destination, ITuple template);
    public ITuple inp(Location current, AgentLocation destination, ITuple template);
    public ITuple[] ing(Location current, AgentLocation destination, ITuple template);
    public ITuple rd(Location current, AgentLocation destination, ITuple template);
    public ITuple rdp(Location current, AgentLocation destination, ITuple template);
    public ITuple[] rdg(Location current, AgentLocation destination, ITuple template);
    public RegisteredReaction[] addStrongReaction(LocalizedReaction[] reactions);
    public RegisteredReaction[] addWeakReaction(Reaction[] reactions);
    public void removeReaction(RegisteredReaction[] reactions);
    public boolean isRegisteredReaction(RegisteredReaction reaction);
    public RegisteredReaction[] getRegisteredReactions();
}

```

Figure 2: The class `LimeTupleSpace`, representing a transiently shared tuple space.

changes in the configuration of the system.

Together, the `LimeSystem` tuple space and the other application-defined transiently shared tuple spaces enable the definition of a fully context aware style of computing.

2.3 Programming with LIME

We complete the presentation of the LIME model by concisely illustrating the application programming interface provided in the current implementation⁴ of LIME.

The class `LimeTupleSpace`, whose public interface is shown⁵ in Figure 2, embodies the concept of a transiently shared tuple space. In the current implementation, agents are single-threaded and only the thread of the agent which creates the tuple space is allowed to perform operations on the `LimeTupleSpace` object; accesses by other threads fail by returning an exception. This represents the constraint that the ITS must be permanently and exclusively attached to the corresponding mobile agent. The name of the tuple space is specified as a parameter of the constructor.

Agents may also have *private* tuple spaces, i.e., not subject to sharing and not appearing in the `LimeSystem` tuple space. A private `LimeTupleSpace` can be used as a stepping stone to a shared data space, allowing the agent to populate it with data prior to making it publicly accessible, or it can be useful as a primitive data structure for local data storage. All tuple spaces are initially created private, and sharing must be explicitly enabled by calling the instance method `setShared`. The method accepts a boolean parameter specifying whether the transition is from private to shared (`true`) or vice versa (`false`). Calling this method effectively triggers engagement or disengagement of the corresponding tuple space. The sharing properties can also be changed in a single atomic step for multiple tuple spaces owned by the same agent by using the static version of `setShared` (see Figure 2). Engagement or disengagement of an entire host, instead, can be triggered explicitly by the programmer by using the methods `engage` and `disengage`, provided by the `LimeServer` class, not shown here. Otherwise, they are implicitly called by the run-time support according to connectivity. The `LimeServer` class is essentially an interface towards the run-time support, and exports additional system-related features, e.g., loading of an agent into a local or remote run-time support, setting of properties, and so on. In particular, it also allows the programmer to define whether transient sharing is constrained to a host-level tuple space, or whether it spans the whole federated tuple space.

`LimeTupleSpace` contains the Linda operations needed to access the tuple space, as well as the operation variants annotated with location parameters. The only requirement for tuple objects is to implement the interface `ITuple`, which is defined in a separate package providing access to a lightweight tuple space implementation. As for location parameters, LIME provides two classes, `AgentLocation` and `HostLocation`, which extend the common superclass `Location`, enabling the definition of globally unique location identifiers for hosts and agents. Objects of these classes are used to specify different scopes for LIME operations, as described earlier. For instance, a probe `inp(cur, dest, t)` may be restricted to the tuple space of a single agent if `cur` is of type `AgentLocation`, or it may

⁴The LIME Web site [34] contains extensive documentation and programming examples.

⁵Exceptions are not shown for the sake of readability.

```

public abstract class Reaction {
    public final static short ONCE;
    public final static short ONCEPERTUPLE;
    public ITuple getTemplate();
    public ReactionListener getListener();
    public short getMode();
    public Location getCurrentLocation();
    public AgentLocation getDestinationLocation();
}
public class UbiquitousReaction extends Reaction {
    public UbiquitousReaction(ITuple template, ReactionListener listener, short mode);
}
public class LocalizedReaction extends Reaction {
    public LocalizedReaction(Location current, AgentLocation destination,
        ITuple template, ReactionListener listener, short mode);
}
public class RegisteredReaction extends Reaction {
    public String getTupleSpaceName();
    public AgentID getSubscriber();
    public boolean isWeakReaction();
}
public class ReactionEvent extends java.util.EventObject {
    public ITuple getEventTuple();
    public RegisteredReaction getReaction();
    public AgentID getSourceAgent();
}
public interface ReactionListener extends java.util.EventListener {
    public void reactsTo(ReactionEvent e);
}

```

Figure 3: The classes `Reaction`, `RegisteredReaction`, `ReactionEvent`, and the interface `ReactionListener`, required for the definition of reactions on the tuple space.

refer the whole host-level tuple space, if `cur` is of type `HostLocation`. The constant `Location.UNSPECIFIED` is used to allow any location parameter to match. Thus, for instance, `in(cur, Location.UNSPECIFIED, t)` returns a tuple contained in the tuple space of `cur`, regardless of its final destination, including also misplaced tuples. Note how typing rules allow the proper constraint of the current and destination location according to the rules of the LIME model. For instance, the `destination` parameter is always an `AgentLocation` object, as agents are the only carriers of “concrete” tuple spaces in LIME. In the current implementation of LIME, probes are always restricted to a local subset of the federated tuple space, as defined by the location parameters. An unconstrained definition, as the one provided for `in` and `rd`, would involve a distributed transaction in order to preserve the semantics of the probe across the federated tuple space.

All the operations retain the same semantics on a private tuple space as on a shared tuple space, except for blocking operations. Since the private tuple space is exclusively associated to one agent, the execution of a blocking operation when no matching tuple is present would suspend the agent forever, effectively waiting for a tuple that no other agent can possibly insert. Hence, blocking operations always generate a run-time exception when invoked on a private tuple space.

The remainder of the interface of `LimeTupleSpace` is devoted to managing reactions; other relevant classes for this task are shown in Figure 3. Reactions can either be of type `LocalizedReaction`, where the current and destination location restrict the scope of the operation, or `UbiquitousReaction`, that specifies the whole federated tuple space as a target for matching. The type of a reaction is used to enforce the proper constraints on the registration through type checking. These two classes share the abstract class `Reaction` as a common ancestor, which defines a number of accessors for the properties established for the reaction at creation time. Creation of a reaction is performed by specifying the template that needs to be matched in the tuple space, a `ReactionListener` object that specifies the actions taken when the reaction fires, and a mode. The `ReactionListener` interface requires the implementation of a single method `reactsTo` that is invoked by the run-time support when the reaction actually fires. This method has access to the information about the reaction carried by the `ReactionEvent` object passed as a parameter to the method. The reaction mode can be either of the constants `ONCE` or `ONCEPERTUPLE`, defined in `Reaction`. Reactions are added to the ITS by calling either `addStrongReaction` or `addWeakReaction`, depending on the desired semantics. As we discussed earlier, in the current implementation strong reactions are confined to a single host, and hence only a `LocalizedReaction` can be passed to the first method. Registration of a reaction returns an object `RegisteredReaction`, that can be used to deregister a reaction with the method `removeReaction`, and provides additional information about the registration process. The decoupling between the reaction used for the registration and the `RegisteredReaction` object returned allows for registration of the same reaction on different

```

System Producer/Consumer
  Program Producer(i) at  $\lambda$ 
    declare
      jobs : TupleSpace
      jobName :  $\{\varepsilon, A, B\}$ 
      jobInfo : String
    initially
      jobs =  $\{\}$   $\parallel$  jobName =  $\varepsilon$ 
    assign
      jobName := A  $\parallel$  jobInfo := newInfo(A)
       $\parallel$  jobName := B  $\parallel$  jobInfo := newInfo(B)
       $\parallel$  (out(jobs, createTuple(jobName, jobInfo))  $\parallel$  jobName =  $\varepsilon$ ) if jobName  $\neq$   $\varepsilon$ 
  end
  Program Consumer(i) at  $\lambda$ 
    declare
      jobs : TupleSpace
      curJob : Tuple
      jobName :  $\{A, B\}$ 
    initially
      jobs =  $\emptyset$   $\parallel$  curJob =  $\varepsilon$ 
    assign
      curJob := in(jobs, createTuple(jobName, String)) if curJob =  $\varepsilon$ 
       $\parallel$  curJob :=  $\varepsilon$   $\parallel$  doJob(curJob) if curJob =  $\varepsilon$ 
  end
Components
  Producer(0)  $\parallel$  Consumer(1)  $\parallel$  Consumer(2)
Interactions
   $\langle \forall a, b, T : a, b \in \mathcal{C}, T \in \mathcal{T} :: a.T \approx b.T \rangle$ 
end

```

Figure 4: A Linda producer/consumer specified in Mobile UNITY.

ITSS and for the same reaction to be registered with strong and, subsequently, with weak semantics.

3 Formal Semantics of LIME

The ultimate goal of our research is rapid development of dependable mobile applications. We contend that precise understanding of the underlying model and its implementation is essential to achieving a high level of dependability. Of course, highly complex formal specifications may be ignored or may actually become a source of confusion for the middleware user. For this reason we opted in favor of modeling the semantics of LIME by using an existing operational model called Mobile UNITY [15, 30], a state-based model consisting of a notation system and associated proof logic, which extends the original UNITY [8] model. Mobile UNITY was specifically designed to express mobility in all its forms and to enable one to reason formally about systems of mobile components, and has been successfully applied to this end [26, 16].

Ultimately, the user of LIME is provided two primary benefits from the formal specification. First, the formalization provides a precise specification of the middleware behavior, aside from the actual implementation. Second, the Mobile UNITY proof logic, in combination with the LIME specification and an application specification allow one to reason formally about the behavior of a complex, mobile application.

The remainder of this section is divided in two parts. In Section 3.1 we present a formal specification of the Linda model which LIME builds upon. Section 3.2 defines the semantics of LIME. Despite reliance on Mobile UNITY, the presentation is self-contained and does not require any prior knowledge of this formal language.

3.1 Formalizing Linda

Figure 4 shows a Mobile UNITY specification⁶ for a Linda-based producer/consumer system where jobs are exchanged through the tuple space. The producer randomly and continuously generates jobs of different names (A and B), putting the job name and its description into a tuple space called *jobs*. Each of the two consumers removes jobs (whose name is randomly determined at initialization time) one at a time from the tuple space, and performs the corresponding actions.

The **Program** sections describe the behavior of each kind of concurrent program⁷. The **Components** section defines the program instances (or agents) that make up the system. Instances of the same program are distinguished by an identifier that is set at creation time, and the statements of the program are essentially duplicated for each component instance. Each program contains a **declare** section for naming variables and declaring their types, and an **initially** section for defining the allowed initial values of the variables. Variables that are not explicitly assigned an initial value in the **initially** section assume an arbitrary value compatible with the type. The **assign** section specifies the guarded assignments that define the state transitions of the system. System execution involves non-deterministically selecting a statement in a weakly fair manner from one of the component programs, and evaluating its guard. If it is true, the statement is executed, otherwise it is skipped. Statements are separated by the \parallel operator. The \parallel operator is used to construct multiple assignments to be executed in a single atomic step. For example, the consumer program contains two statements; the first removes a job tuple from the tuple space, while the second simulates the performance of a job through a function `doJob`, whose details are not relevant here, and resets the value of *curJob* to ε in the same atomic step. The guard in the first statement prevents the consumer from taking a new job before the previous one is completed (i.e., before *curJob* is reset to ε).

In Mobile UNITY, and different from the original UNITY, variables with the same name in different programs are distinct. For example, the tuple space variable *jobs* in the producer is distinct from the one in either of the consumers. Nevertheless, in Linda these tuple spaces should always be the same. To accomplish this in Mobile UNITY, we exploit the sharing construct \approx , which allows one to express symmetric and transitive sharing between variables belonging to different programs. Let us assume that \mathcal{C} is the set of all component names and \mathcal{T} is the set of names of all the variables of type TupleSpace. In our example, $\mathcal{C} = \{Producer(0), Consumer(1), Consumer(2)\}$, and $\mathcal{T} = \{jobs\}$.

Hence, by stating⁸ in the **Interactions** section that:

$$\langle \forall a, b, T : a, b \in \mathcal{C}, T \in \mathcal{T} :: a.T \approx b.T \rangle$$

we force any tuple spaces with the same variable name to be shared, even if they are declared within the name spaces of different agents. In our example, the *jobs* tuple spaces of the producer and consumers become shared by virtue of this statement; hence, when the producer writes a tuple to its tuple space it is immediately available in the consumers' tuple space. Moreover, in this case sharing is unconditional, that is, the tuple spaces are *always* shared. When describing LIME, we will remove this assumption and specify the conditions under which transient sharing is enabled, by exploiting more sophisticated forms of the sharing construct.

Placement of the above statement in the **Interactions** section is intentional. This section of the specification is reserved in Mobile UNITY for statements that involve more than one component. The formal semantics of Linda, and next of LIME, is hence constituted by two parts. The first one are the statements we place in the **Interactions** section, which effectively capture a significant part of the run-time communication among programs, as enabled by the middleware. The **Interactions** section of *every* specification exploiting the LIME model must contain the **Interactions** statements we describe in the LIME formal semantics. The second part is the definition of the various constructs, such as the Linda operations *in* and *out* appearing in the example, that are available to the programmer. These constructs are technically defined as *macros*, whose meaning is represented in terms of the basic Mobile UNITY statements. The examples we use here for Linda and next for LIME provide the reader with the opportunity to see how these two constituents of our formal semantics are exploited in the specification of an application.

⁶In the following, we use different fonts for improving the readability of our semantics, by distinguishing between Mobile UNITY constructs (e.g., **reacts-to**), constructs that are available to the specifier as part of LIME (e.g., *in*), auxiliary macros used for structuring the specification but not available to the programmer (e.g., *copy*), types (e.g., Tuple), constants and labels (e.g., CUR), and normal variable names (e.g., *events*).

⁷For now we will ignore the λ variable on each program.

⁸This statement uses a UNITY-like three part notation of the form $\langle \text{op } \textit{quantified Variables} : \textit{range} :: \textit{expression} \rangle$. The variables from *quantifiedVariables* take on all possible values permitted by *range*. If *range* is missing, the first colon is omitted and the domain of the variables is restricted by context. Each such instantiation of the variables is substituted in *expression* producing a multiset of values to which **op** is applied, yielding the value of the three-part expression. If no instantiation of the variables satisfies *range*, the value of the three-part expression is the identity element for **op**, e.g., *true* when **op** is \forall .

In this section, we begin the formal description of Linda by focusing first on the essential data structures, then on the operations.

Data Structures. The fundamental data structures used in Linda are tuples and tuple spaces. Tuples are represented by the type `Tuple`, instances of which are generated by using the `createTuple` function. Parameters to this function can include actual values, as in the producer of Figure 4, or any combination of actuals and formals (types), as in the consumer. The result returned by this function is a tuple that is augmented by some fields that were not present in the tuple originally written by the specifier, but that are necessary to the semantics of the operations. In formalizing Linda, the only field added by `createTuple` is a unique identifier; when formalizing LIME in the next section, `createTuple` is modified to add also fields for the current and destination location. These fields remain part of the tuple, accessible to the operations defined in our formalization, but not accessible to the specifier.

Although we do not provide a formal definition of `createTuple`, as it is tedious and not interesting, we define here a general notation for setting and accessing these added tuple fields. The tuple identifier is associated to a *field label* `ID`; later on, we will use the labels `CUR` and `DEST` for the current and destination location fields, respectively. These field labels are used to create template expressions. For instance, $[ID: ID] \oplus t$ sets the identifier field of tuple t to the formal `ID`, while $[ID: 42] \oplus t$ sets it to the actual `42`. Moreover, field labels are also used to retrieve the field value. Hence, $t.ID$ returns the current location of the tuple or, in conformance with the semantics of Linda, undefined if the field is currently set to a formal.

Having augmented each tuple with an extra identifier field allows us to formalize the tuple container as a set of type `TupleSpace` (as opposed to a multiset), as long as the identifier value is set properly. This definition presents the specifier with constructs that allow for multiple tuples with identical data to exist in the tuple space, as in Linda, but allows us to exploit set operations in the specification of the underlying semantics.

Constructs. The final statement of the producer program demonstrates the use of the operation $out(T, t)$ to insert a tuple into the tuple space. Because a process can access multiple tuple spaces, the tuple space variable appears as a parameter of the operation. The formal definition describes the tuple space change resulting from the execution of the `out` operation, namely the insertion of the tuple t :

$$out(T, t) \triangleq T := T \cup \{[ID: newId()] \oplus t\}$$

In addition to inserting the tuple into the tuple space, the uniqueness of the tuple is established by setting the identifier field to a system-wide unique tuple identifier, returned by the function `newId()`. It should also be noted that by the assignment semantics of Mobile UNITY, the value in braces is actually the value of the tuple, and thus a *copy* of the tuple is effectively made and inserted into the tuple space.

Before turning to the formal definitions of the `in` and `rd` operations for retrieving data from the tuple space, we introduce a few auxiliary definitions. First, to formally express the fact that a tuple θ matches a template p , we use the notation $\mathcal{M}(\theta, p)$. We do not provide a fully formal definition for the semantics of matching since it is not fundamental to our model. Moreover, we extend the definition of formal fields to allow the specification of subtypes. For example, $\langle \text{"foo"}, \text{Integer } i : 1 \leq i \leq 10 \rangle$ requires a matching tuple to have, in its second field, an integer value between 1 and 10. While we do not use this functionality directly in our example program, it is used in the definitions of several LIME constructs in the next subsection⁹.

Next we define a predicate that identifies whether a tuple matching a pattern exists in a given tuple space, and two macros that identify a matching tuple in a given tuple space and either remove or copy it:

$$\begin{aligned} \text{matchExists}(T, p) &\equiv \langle \exists \theta : \theta \in T :: \mathcal{M}(\theta, [ID: ID] \oplus p) \rangle \\ t := \text{remove}(T, p) &\triangleq \langle \parallel \theta : \theta = \theta'. (\mathcal{M}(\theta', [ID: ID] \oplus p) \wedge \theta' \in T) :: T, t := T - \{\theta\}, \theta \rangle \\ t := \text{copy}(T, p) &\triangleq \langle \parallel \theta : \theta = \theta'. (\mathcal{M}(\theta', [ID: ID] \oplus p) \wedge \theta' \in T) :: t := \theta \rangle \end{aligned}$$

In both the `remove` and `copy` macros, the value of the returned tuple is bound to the variable t , leaving the pattern p with its original value. These macros model non-deterministic tuple selection by means of *non-deterministic*

⁹Moreover, while this feature is used in the formal semantics, matching on subtypes (e.g., using the inheritance relationship) is currently not supported by the middleware.

assignment [1]. In a non-deterministic assignment of the form $x := x'.Q$, the variable x is assigned a value x' , selected non-deterministically among those satisfying condition Q . In our case, we use a similar notation to select non-deterministically a single matching tuple θ' , bind it to the tuple θ , and use it to quantify the three-part notation. Because a single tuple is selected and bound, the parallel operator in the three-part notation serves the purpose of creating a quantified statement, and does not entail parallel execution of multiple statements.

With these helper functions, the definitions of the `rd` and `in` constructs follow naturally:

$$\begin{aligned} t := \mathbf{in}(T, p) &\triangleq t := \mathbf{remove}(T, p) \quad \mathbf{if} \ \mathbf{matchExists}(T, p) \\ t := \mathbf{rd}(T, p) &\triangleq t := \mathbf{copy}(T, p) \quad \mathbf{if} \ \mathbf{matchExists}(T, p) \end{aligned}$$

The previous definitions describe the *blocking* forms of the `in` and `rd` operations. In Linda, a process that encounters a blocking operation suspends itself until a matching tuple is found. Instead, in Mobile UNITY there is no direct notion of process blocking: statements are selected non-deterministically. Nevertheless, if a statement is selected when no matching tuple exists in the tuple space, it is equivalent to a skip. Thus, it is as if the statement to remove the tuple was *blocked* waiting for a matching tuple. In the example, when no tasks exist in the tuple space the consumer is effectively blocked waiting for the `in` operation to become enabled. With this in mind, it is possible and meaningful to put these blocking operations in parallel with other statements. For example:

$$t := \mathbf{in}(T, p) \parallel \mathit{count} := \mathit{count} + 1$$

is expected to count the number of tuples removed. However, after the macro expansion of the `in` operation, the semantics of this statement are such that each time it is selected for execution, the counter increments even if no matching tuple is taken from the tuple space. This is because only the assignment to the left of the parallel bar is inhibited until a match is found. To allow the more meaningful style of parallel assignment in which both assignments are inhibited until a match is found, we expose the `matchExists` predicate, enabling the following assignment statement with the correct counting semantics:

$$(t := \mathbf{in}(T, p) \parallel \mathit{count} := \mathit{count} + 1) \ \mathbf{if} \ \mathbf{matchExists}(T, p)$$

It should be noted that the same parameters, namely (T, p) , must be used for the `in` operation and the `matchExists` to have the desired effect.

Additional constructs. While the above constructs define the basic Linda operations, there are a number of extensions that have been shown to be useful when programming with Linda, some of which we exploit in LIME. In contrast to the blocking version of the `in` and `rd` the non-blocking operations should return ε if no matching tuple exists in the tuple space at the moment when the statement is selected. Their formal definition builds upon the blocking operations:

$$\begin{aligned} t := \mathbf{inp}(T, p) &\triangleq \begin{array}{l} t := \mathbf{remove}(T, p) \quad \mathbf{if} \ \mathbf{matchExists}(T, p) \\ \parallel t := \varepsilon \quad \mathbf{if} \ \neg \mathbf{matchExists}(T, p) \end{array} \end{aligned}$$

The definition of the probing read `rdp` is identical to `inp` except the `copy` macro is used in place of `remove`, leaving the tuple space unchanged after the execution of the operation involving a successful match. The group operations that remove or copy *all* matching tuples (e.g., `ing` and `rdg`) can be formalized in a similar manner.

3.2 Formalizing LIME

In this section, we build upon the formalization presented thus far, and extend it to encompass transiently shared tuple spaces and the other constructs that are peculiar to LIME. The constructs defined in the previous section, namely `out`, `in`, `rd`, `inp`, and `rdp`, are used in the formalization of the LIME constructs, but this section provides new semantics that deal with the current and destination location fields. The functions `createTuple` and `matchExists` are modified respectively to include the creation of the `CUR` and `DEST` fields and to match tuples using these additional fields. While the overall semantics remain the same, operations work with tuple spaces that are transiently, rather than permanently, shared.

```

System Producer/Consumer
Program Producer(i) at  $\lambda$ 
  declare
    jobs : TupleSpace
    jobName : { $\varepsilon$ , A, B}
    jobInfo : String
  initially
    jobs = {} || jobName =  $\varepsilon$ 
  assign
    jobName := A || jobInfo := newInfo(A)
    || jobName := B || jobInfo := newInfo(B)
    || (out(jobs, createTuple(jobName, jobInfo)) || jobName =  $\varepsilon$ )   if jobName  $\neq$   $\varepsilon$ 
    || out[Consumer(1)](jobs, createTuple(jobName, jobInfo)) || jobName =  $\varepsilon$ )   if jobName  $\neq$   $\varepsilon$ 
    || out(jobs, createTuple(PRIORITY, newInfo(PRIORITY)))
  end
Program Consumer(i) at  $\lambda$ 
  declare
    jobs : TupleSpace
    curJob, priorityTemplate : Tuple
    jobName : {A, B}
  initially
    curJob =  $\varepsilon$  || priorityTemplate = createTuple(PRIORITY, String) || isEnabled(pri)
  assign
    curJob := in(jobs, createTuple(jobName, String))   if curJob =  $\varepsilon$ 
    || curJob :=  $\varepsilon$  || doJob(curJob)
    || pri :: doJob( $\tau$ ) weakReaction[AgentID, AgentID](jobs, priorityTemplate, ONCEPERTUPLE)
  end
Components
  Producer(0) || Consumer(1) || Consumer(2)
Interactions
  ;; Transient sharing dependent on connectivity
   $\langle \forall a, b, c, T : a, b, c \in \mathcal{C} \wedge T \in \mathcal{T} :: a.T \approx b.T$    when  $a \kappa b$ 
  engage  $a.T \cup b.T$ 
  disengage  $(\cup c : c \kappa a :: c.T \downarrow c), (\cup c : c \kappa b :: c.T \downarrow c)$ 

  ;; Migration of misplaced tuples
  ||  $\langle || \theta, a, b, T : a, b \in \mathcal{C} \wedge T \in \mathcal{T} \wedge \theta = \theta'.(\theta' \in b.T \wedge \theta'.CUR = a \wedge \theta'.DEST = b) :: b.T := b.T - \{\theta\} + \{[CUR: b] \oplus \theta\} \rangle$  reacts-to true
end

```

Figure 5: A LIME producer/consumer in Mobile UNITY.

Components and connectivity. In a Mobile UNITY specification, as shown in Figure 5, a program is the unit of execution and migration. In our LIME specification, a program represents the specification of the behavior of a mobile agent. In Mobile UNITY, each program component has a special location variable λ , which we use to identify the host where the mobile agent is executing. The structure of this location variable is completely application-dependent, and for our purposes it can safely be assumed to contain the IP address or the symbolic name of a host. Agents migrate by assigning a new value to this location variable, e.g., $\lambda := \text{lime.sf.net}$. We formally express connectivity among agents as a symmetric and transitive relation κ . Two agents a and b are *connected* when they are on the same host, when the hosts they reside on are directly connected, or when the hosts they reside on are transitively connected. This is expressed by the predicate $a \kappa b$. When an agent migrates, the κ relation changes to reflect the new configuration. Changes in connectivity among hosts are intentionally left outside of the specification, but nevertheless contribute to the κ relation. This is merely a consequence of choosing agents as the only active components in the system.

In applications involving physical and/or logical mobility, it is often necessary to gain information about the current connectivity of the mobile components. The natural course of action to support this need during the specification phase would be to expose the κ relation to the specifier. However, this would actually grant the specifier more power than can actually be realized in the mobile environment. For instance, consider¹⁰ the case where an agent a queries

¹⁰Here and in the following we use agent names like a also to refer to the corresponding agent identifier. While this constitutes a stretch of the notation, it greatly simplifies formulas. Moreover, the context in which the variable is used is sufficient to remove ambiguity.

whether agent b is connected to agent c by directly using the predicate $b \kappa c$. If b and c are out of range of a , there is no reasonable way for a to resolve the query about components it cannot communicate with. However, it is both useful and reasonable to allow queries about a 's current connectivity state, e.g., $a \kappa b$. In our model, the latter predicate can be expressed, from the perspective of agent a , as `isConnected(b)`. The use of `isConnected` effectively restricts access to the κ relation, limiting its range only to a 's current connectivity context. Based on this function and the other LIME constructs described in the following, the formal semantics of the LimeSystem tuple space can be modeled straightforwardly.

Transient sharing of tuple spaces. In LIME, as opposed to Linda, identically-named tuple space variables defined in different agents must be shared transiently, *only when connectivity is available*. This is expressed in our formalization by using the sharing construct \approx , as in the formalization of Linda. The full form of this construct allows one to specify the condition enabling sharing (**when**), the value the shared variable assumes when this condition is established (**engage**), and the value each variable takes when the condition is falsified (**disengage**). In our case, the condition for sharing is the existence of connectivity, the engagement value is the union of the content of the tuple spaces, and the disengagement values partition the content of the shared tuple space according to tuple location and the new connectivity.

For notational convenience we define an operator, \downarrow , to generate a set containing all tuples from a tuple space T whose current location field refers to an agent a . Intuitively, $T \downarrow a$ selects all tuples for which a is responsible:

$$T \downarrow a \triangleq \langle \mathbf{set} \theta : \theta \in T \wedge \theta.CUR = a :: \theta \rangle$$

where `CUR` is the field label needed to access the current location field, with a notation similar to the one we used in the previous section for the tuple identifier. With this definition, transient sharing can be expressed by substituting the unconditional sharing clause in the **Interactions** section of the Linda semantics, with the following:

$$\langle \forall a, b, c, T : a, b, c \in \mathcal{C} \wedge T \in \mathcal{T} :: a.T \approx b.T \quad \begin{array}{l} \mathbf{when} \ a \kappa b \\ \mathbf{engage} \ a.T \cup b.T \\ \mathbf{disengage} \ \langle \cup c : c \kappa a :: c.T \downarrow c \rangle, \langle \cup c : c \kappa b :: c.T \downarrow c \rangle \end{array} \rangle$$

where \mathcal{C} and \mathcal{T} have the meaning described earlier. According to this expression, the two tuple space variables named T belonging to agents a and b become shared when connectivity is established between the agents, i.e., when $a \kappa b$ is true. The shared variable T is assigned the set union of the content of the tuple spaces. Correctness of the engagement clause relies on the commutativity of the union operator to generate the same value whether $a.T \approx b.T$ or $b.T \approx a.T$, and on the fact that the tuple space is a set (rather than a multiset), thus ensuring that no tuple duplication takes place when $a.T \approx a.T$. On disengagement, i.e., when $a \kappa b$ is falsified, the content of the shared tuple space variable is partitioned such that the contents of the tuple spaces of each agent reflect the tuples available under the new connectivity. In other words, after a disconnection the visible tuples are only those whose responsible agents are still connected to one another. In our example, the result of transient sharing is that either of the consumers can remove jobs created by the producer only when the two components are connected.

Managing misplaced tuples. In Figure 5, the `out` operation is used with a location parameter to specify that a job tuple should be migrated to the tuple space of *Consumer(1)*. While this does not guarantee that *Consumer(1)* will process the tuple, the migration makes it possible for *Consumer(1)* to process the tuple even when not connected with the producer. The semantics of this extended operation can be modeled by the following:

$$\mathbf{out}[d](T, t) \triangleq \mathbf{out}(T, [\mathbf{CUR}: a, \mathbf{DEST}: d] \oplus t)$$

which stores the tuple t with a (the agent that executed the operation) as the current location and d as the destination. In this definition we rely on the definition of `out` given earlier for Linda, but we modify the tuple parameter by extending it with location fields. The setting of the identifier field and the change in the tuple space content is performed by the `out` we described in Section 3.1, while the values of the `CUR` and `DEST` fields are set explicitly. Note that this relies on the aforementioned redefinition of the `createTuple` function to ensure the existence of these two fields. Finally, for notational convenience we allow the specifier to leave out the values for some or all locations, as shown in Table 1 of Section 2.2. When missing, they are assumed to be the identifier of the agent performing the `out`.

Since this statement is not guarded, it is executed independently of whether the destination d is connected to a or not. In the case where the writing and the destination agents are connected when the `out` is issued, tuple

migration is immediate. In the case where connectivity does not immediately exist, the migration occurs as soon as connectivity becomes available. This opportunistic action is modeled by an additional statement that employs the reactive construct **reacts-to** provided by Mobile UNITY.

Reactive statements extend Mobile UNITY statements with the capability to specify actions that must be executed immediately after a given condition is established rather than eventually as dictated by the fair interleaving semantics. Reactive statements can appear within an individual program specification or within the **Interactions** section. The reactive statements of the entire system are logically combined to form a single reactive program that is executed to fixed point after the execution of every conventional (non-reactive) statement. It is the responsibility of the specifier to ensure fixed point of the reactive program is actually reached¹¹. In our model this condition is guaranteed by construction, except for the case when a reaction generates a tuple that triggers another reaction—a condition that is nonetheless easy to analyze at the application level.

For specifying tuple migration, we define the following reactive statement in the **Interactions** section:

$$\langle \langle \theta, a, b, T : a, b \in \mathcal{C} \wedge T \in \mathcal{T} \wedge \theta = \theta'. (\theta' \in b.T \wedge \theta'. \text{CUR} = a \wedge \theta'. \text{DEST} = b) :: b.T := b.T - \{\theta\} + \{[\text{CUR}: b] \oplus \theta\} \rangle \rangle$$

reacts-to *true*

Although connectivity is not explicitly mentioned in this formula, it is implicit that tuples can migrate from a to b only when connectivity is available. The interpretation of the above formula from the perspective of agent b shows that b selects a misplaced tuple (with current location a and destination b) from its own tuple space, $b.T$. Since $b.T$ is shared with the other tuple spaces, the presence of a tuple belonging to a inside b 's tuple space guarantees that a is connected. Stated differently, a tuple cannot exist in a tuple space if the agent specified in the tuple's current location field is not present (i.e., connected). In the formula above, the misplaced tuple is then “migrated” by removing it from the tuple space, and reinserting it with the CUR field properly reset to agent b .

The semantics of reactions guarantee that tuples migrate immediately when connectivity is available. If a tuple is written while the current and destination agents are connected, the reaction fires immediately after the out operation. Alternately, when components holding misplaced tuples come into range, the reaction above executes in the same atomic step that engages the tuple space variables. In both cases, agents perceive an instantaneous tuple migration to the destination agent.

Formalizing location-extended query constructs. The location-extended variants of the other operations can be modeled by using the earlier defined Linda operations along with the CUR and DEST tuple fields. Here, we focus on the formalization of in:

$$t := \text{in}[a, b](T, p) \triangleq t := \text{in}(T, [\text{CUR}: b, \text{DEST}: c] \oplus p)$$

The formalizations of inp, ing, and the rd variants directly follow. In the above formula, the values in square brackets are the current and destination location parameters, and any of the tuple space projections defined in Section 2 can be defined easily. Agents can be named explicitly using their identifier; the formal AgentID is used to match any agent. To express queries ranging over an entire host, we use the subtype matching definition described in Section 3.1. For example, the query necessary to remove a tuple destined to agent b from the host-level tuple space named $jobs$ of the co-located agent a is:

$$\text{in}[\langle \text{AgentID } i : i.\lambda = a.\lambda \rangle, b](jobs, p)$$

Care must be taken to ensure that the predicate is computable given the connectivity constraints of the system. For example, it is not reasonable to write a predicate that specifies for the DEST field all agents residing on a host that is not accessible. Both predicates used here are computable, and are available in the implementation of LIME. For notational convenience, we allow the specifier to use a non-augmented version of these operations, with the default meaning to query the current federated tuple space. As discussed in Section 2.3, the CUR field of the probe and group operations should be restricted to a single host, although this is not enforced by the formalization.

¹¹The introduction of reactive statements is reflected in the Mobile UNITY logic, and proper inference rules are needed to prove the correctness of reactive programs. The interested reader can find details about this in [15].

Strong and weak reactions. The example in Figure 5 extends the producer and the consumers to handle *high priority jobs*. These are generated by the producer and should be performed by any of the consumers “as soon as possible”, i.e., as allowed by connectivity. If the producer and a consumer are connected when the job is generated, the consumer should immediately remove the corresponding tuple and perform the job. Vice versa, the consumer should take and perform the job as soon as connectivity is restored. From the producer’s perspective, the only change is the addition of a statement (the last one) that outputs a job tuple with `PRIORITY` as the job name. The consumer program, on the other hand, must be able to *react* to the presence of a high priority tuple. LIME reactions provide the natural solution for implementing this behavior. As discussed in Section 2, LIME provides two kinds of reactions: strong reactions fire in the same atomic step detecting the presence of a matching tuple, while weak reactions fire only eventually after detection. In the example, the consumer program is augmented with a statement (the last one) containing a weak reaction that is fired whenever a tuple containing high priority jobs appears in the shared tuple space, and executes the corresponding job. Using a weak reaction in the example makes it closer to the specification of a real application designed using the current LIME implementation where, as we mentioned in Section 2, performance considerations driven by the current target scenario constrain strong reactions to fire only upon detection of a local tuple. In the formalization that follows, however, we do not model explicitly the constraints the current middleware imposes on strong reactions, to keep the formal semantics of our model as free as possible from the implementation concerns specific to a given application scenario.

We model LIME reactions as:

$$\begin{aligned} \rho &:: s(\tau) \text{strongReaction}[x, y](T, p, mode) \\ \rho &:: s(\tau) \text{weakReaction}[x, y](T, p, mode) \end{aligned}$$

where s is the statement to be performed when a tuple matching pattern p is found in the tuple space T . As in the query operations, $[x, y]$ indicates a projection of the tuple space over which to evaluate the reaction. The variable τ is a free variable that can appear within the statement s and is bound to the matched tuple when the action fires. The variable $mode$ can assume either of the `ONCE` or `ONCEPERTUPLE` values. The label ρ uniquely identifies the reaction and can be used as a parameter for the registering and deregistering operations, `enableReaction(ρ)` and `disableReaction(ρ)`.

The formalization of reactions relies on the **reacts-to** construct provided by Mobile UNITY. Nevertheless, some additional bookkeeping is necessary to deal with explicit registration and deregistration of reactions (which is not provided in Mobile UNITY) and to ensure the proper execution pattern with respect to the reaction mode. The details of the formalization are discussed below for strong reactions:

$$\begin{aligned} \rho &:: s(\tau) \text{strongReaction}[x, y](T, p, mode) \triangleq \\ &\langle \parallel \theta : \theta = \theta'. (\mathcal{M}(\theta', [\text{ID: ID, CUR: } x, \text{ DEST: } y] \oplus p) \wedge \theta' \in T \\ &\quad \wedge (mode = \text{ONCE} \Rightarrow \text{reactedSet}_\rho = \{\}) \\ &\quad \wedge (mode = \text{ONCEPERTUPLE} \Rightarrow \theta'.\text{ID} \notin \text{reactedSet}_\rho)) \\ &\quad :: s\{\tau/\theta\} \parallel \text{reactedSet}_\rho := \text{reactedSet}_\rho \cup \{\theta.\text{ID}\} \rangle \\ &\quad \text{reacts-to} \\ &\quad \text{matchExists}(T, [\text{ID: ID, CUR: } x, \text{ DEST: } y] \oplus p) \wedge \text{enabled}_\rho \end{aligned}$$

To keep track of the tuples that have been reacted to and of whether the user has enabled or disabled the reactions, two auxiliary variables, namely reactedSet_ρ and enabled_ρ , are introduced. By subscripting these variables with the unique reaction identifier ρ , we are guaranteed that reactions do not interfere with one another.

The selection of a matching tuple uses the same non-deterministic selection notation as described in the Linda formalization for the `copy` function. The necessity to bind the matched tuple to the free variable τ prohibits the use of the `copy` function here, but the semantics are the same: a single matching tuple is selected from the shared tuple space. The condition for actually firing the reaction (i.e., executing the user action s) depends on the contents of reactedSet_ρ . This variable tracks the identity of the tuples that have already been reacted to. A `ONCE` reaction will only fire if no tuples have been reacted to, i.e., the set is empty. A `ONCEPERTUPLE` reaction will only fire if the tuple chosen by the non-deterministic selection has not been recorded in reactedSet_ρ . The user action $s(\tau)$ can be any non-reactive Mobile UNITY statement¹².

¹²More precisely, transactions cannot appear in user actions either. Since transactions are not used in our formalization of LIME, we redirect the reader interested in the semantic details of Mobile UNITY to [15].

In order to guarantee that at least one tuple can be selected from the tuple space, we include the `matchExists` function as a condition to the reaction. Without this, the θ' variable could not be bound and the formalization would be incorrect. The second condition for the reaction, $enabled_\rho$, tracks whether the user has explicitly enabled or disabled the reaction with the following macros:

$$\begin{aligned} \mathbf{enableReaction}(\rho) &\triangleq (enabled_\rho := true \parallel reactedSet_\rho := \{\}) \quad \mathbf{if} \neg enabled_\rho \\ \mathbf{disableReaction}(\rho) &\triangleq enabled_\rho := false \end{aligned}$$

These functions allow the specifier to control when reactions are able to fire during the lifetime of the system. By clearing $reactedSet_\rho$ when a reaction is re-enabled, it is treated as a new reaction that has not fired on any tuples. This is true for both `ONCE` and `ONCEPERTUPLE`; neither type of reaction retains any *memory* of previous periods when it was enabled and they may react again to the same tuples, if re-enabled. It is also important to remember that the reactive program is executed after every regular statement, including a statement that enables a reaction. This means that a reaction may fire in the same atomic step as the statement that enables it, assuming a matching tuple exists in the tuple space. As mentioned in Section 2, this is a significant departure from more conventional event-based systems.

For proper initialization, each reaction must be either enabled with an empty $reactedSet_\rho$, or disabled. This is indicated in the **initially** section of a program for reaction ρ by using either of the following predicates:

$$\begin{aligned} \mathbf{isEnabled}(\rho) &\triangleq enabled_\rho = true \wedge reactedSet_\rho = \{\} \\ \mathbf{isDisabled}(\rho) &\triangleq enabled_\rho = false \end{aligned}$$

As mentioned earlier, the formalization just introduced does not limit the scope of strong reactions to the local host, as done by our current implementation. This constraint is easy to capture, but we preferred to define here the formal semantics of the LIME model in the most general way.

Moreover, our model already provides, through weak reactions, an alternative reaction semantics that imposes fewer implementation requirements. Weak reactions remove the atomicity between the detection of the matching tuple and the firing of the reaction, thus allowing other agent operations to occur in between the two, and enjoy a scope that can be as large as the whole federated tuple space. We take advantage of the similarities of weak and strong reactions, and actually rely on the latter in the specification of the formal semantics. Specifically, a strong reaction is exploited to detect the appearance of a matching tuple. The user action is instead an asynchronous statement, that hence breaks the atomicity and guarantees only eventual completion of the weak reaction:

$$\begin{aligned} \rho :: s(\tau) \mathbf{weakReaction}[x, y](T, p, mode) &\triangleq \rho :: events_\rho := events_\rho \cup \{\tau\} \mathbf{strongReaction}[x, y](T, p, mode) \\ &\parallel \langle \parallel \theta : \theta = \theta'.(\theta' \in events_\rho) :: events_\rho := events_\rho - \{\theta\} \parallel s\{\tau/\theta\} \rangle \end{aligned}$$

To model the transfer of a matching tuple from a given agent to the subscriber, we rely on the set $events_\rho$. This variable is a temporary holding place for all tuples that should be reacted to, but for whom the user reaction has not yet fired. In some sense, even though $events_\rho$ is local to the agent that registered the reaction, it is analogous to a communication buffer between the agents. This set is populated by the strong reaction in the first statement. The second statement, separated by the \parallel operator, executes asynchronously with respect to the reaction by removing an element from $events_\rho$ and firing the corresponding user action using the matched tuple. As with strong reactions, only tuples in the currently shared tuple space will be reacted to, meaning that any tuples written while hosts are disconnected cannot be reacted to, even in the weak model. Additionally, we must redefine `isEnabled` for weak reactions in order to clear the set $events_\rho$:

$$\mathbf{isEnabled}(\rho) \triangleq enabled_\rho = true \wedge reactedSet_\rho = \{\} \wedge events_\rho = \{\}$$

One point to notice about the definition of `weakReaction` is the use of non-deterministic selection to remove an element from $events_\rho$. Similar to the definition of `remove` in Section 3.1, exactly one tuple will be selected, ensuring that only one user action s will fire at a time. Also, this non-deterministic selection does not guarantee an order in the selection of elements from $events_\rho$, nor does it guarantee that an element will *ever* be selected. This is the weakest constraint we can set in the formalization, and leaves room for stronger guarantees to be enforced by specific implementations. For instance, our current implementation maintains ordering and guarantees selection.

4 Design and Implementation of LIME

In this section we look behind the scenes of the LIME programmer interface, and describe some aspects of the architecture of the middleware.

Tuple Space. One of the early decisions in the design of LIME addressed the implementation of the underlying tuple space. Analysis of available systems revealed that they provide a rich set of features with large variations in terms of expressiveness, performance, and often semantics. The need for a simple, lightweight implementation, combined with the desire to provide support and interoperability with industrial-strength products, led us to the development of an adaptation layer that hides from the rest of the LIME implementation the nature of the underlying tuple space engine. This layer is provided by a separate package called LIGHTS [21], developed by one of the authors, which also includes a lightweight tuple space implementation. In this package, the interface `ITupleSpace` provides access to the core tuple space functionality. Adapter classes implementing this and other interfaces are loaded at startup to translate operations into those of the supported tuple space engines. Currently, adapters exist for LIGHTS's built-in lightweight tuple space and for IBM's TSpaces [12].

In LIME, every `LimeTupleSpace` object contains an instance of the LIGHTS tuple space, accessible through the `ITupleSpace` interface. Instances of `LimeTupleSpace` are created as private, so that only the agent creating the tuple space has access to the contents of the tuple space. Operations on the `LimeTupleSpace` are delegated to the underlying `ITupleSpace` object if the tuple space is currently private. Otherwise, they are delegated to a component of the run-time that enforces the semantics of transient sharing, as described in the following.

Location Parameters. Although locations are not immediately useful in a private tuple space, `LimeTupleSpace` is already equipped with the mechanisms needed to deal with location attributes, since they become relevant when the tuple space is toggled from private to shared. Such mechanisms simply entail the management of two fields that are dynamically added and removed from tuples, and that correspond to the current and destination locations. Upon insertion in the tuple space, the original tuple provided by the programmer is always augmented with these two fields, which are set to the proper value—those specified by the programmer if the operation contains location parameters, conventional default values otherwise. These fields are then stripped off when the tuple is returned as the result of an operation accessing the tuple space, such as `rd`. Although these fields are exploited by the run-time support to deal with locations, they are hidden from the programmer, thus preventing the introduction of inconsistencies by directly altering their value.

Host-Level Sharing. When a private `LimeTupleSpace` is set to shared, a host-level tuple space is created through transient sharing, according to the semantics we presented earlier. The enforcement of these semantics, and in particular of the engagement and disengagement of local tuple spaces, cannot be managed by a single `LimeTupleSpace`, as it requires host-wide management of tuple space access. This management is provided by instances of the class `LimeTSMgr`. On each host, a `LimeTSMgr` object exists per tuple space name: each of these objects is created when a local `LimeTupleSpace` with the corresponding name becomes shared for the first time. Subsequent engagements of local tuple spaces with the same name exploit the same `LimeTSMgr`.

When the `LimeTupleSpace` becomes shared, it surrenders the control of its tuple space to the appropriate `LimeTSMgr`. In practice, the methods of the `LimeTupleSpace` no longer operate directly on the `ITupleSpace`, rather they delegate tuple space operations to the `LimeTSMgr`, which performs the appropriate actions according to the semantics of transient sharing before returning control to the method body and hence to the calling agent. Operation requests are queued and serially executed at the `LimeTSMgr`, which runs in a separate thread of control. This way, synchronization among concurrent accesses performed through different `LimeTupleSpace` instances is obtained structurally, by confining all tuple space accesses to a synchronized queue.

To perform its actions, the `LimeTSMgr` must obtain access to the content of the tuple space being shared. This access can be granted in at least two ways. The first one consists of providing the `LimeTSMgr` with a direct reference to the `ITupleSpace` object holding such content. The other solution consists of moving the data tuples into a single `ITupleSpace` that contains all the shared tuples and is associated with the `LimeTSMgr`¹³, and moving them back to the original `ITupleSpace` upon disengagement. The first approach is more convenient in very dynamic scenarios, where the overhead of moving the tuples back and forth as a consequence of a reconfiguration becomes significant.

¹³In this case, the tuple space associated with the `LimeTSMgr` can be regarded as a concrete representation of the host-level tuple space.

Instead, the second approach, which is chosen by our current implementation, opens up opportunities for optimizing query execution since, rather than searching several tuple spaces, a single one can be searched.

Reactions. In addition to the `ITupleSpace` object, each `LimeTupleSpace` and `LimeTSMgr` contain also a `Reactor` object that is responsible for managing the (de)registration of reactions and the execution of the reactive program. `Reactor`, `LimeTupleSpace`, and `LimeTSMgr` are properly synchronized so that only tuple space operations issued from within the statements of the reactive program are allowed to execute while all others are blocked during the execution of a reaction.

When a strong reaction is registered by the user, a `RegisteredReaction` object is generated and kept in a list held by the local `Reactor`. This list effectively contains the reactive program, which must execute after each tuple space operation that adds tuples to the tuple space. Execution of the reactive program proceeds by iterating through the list of registered reactions. If a tuple matching a reaction is found, the reaction is fired by executing the corresponding `ReactionListener`. Iteration continues until one pass completes with no reactions firing.

Weak reactions are built on top of strong reactions. When a weak reaction is registered, a message containing the reaction template and a reaction identifier is sent to all the hosts involved, based on the location parameters associated with the reaction. On each host, the message causes the registration of a system-defined strong reaction watching for tuples matching the specified template. A similar reaction is registered also locally. Moreover, the `ReactionListener` provided by the programmer is stored along with the reaction identifier within a table associated to the local `LimeTSMgr`. When one of the strong reactions fires on one of the hosts, the corresponding `ReactionEvent` is sent back to the subscribing host, along with information about the reaction identifier. The latter is used to retrieve the correct listener from the table and to execute it by passing the `ReactionEvent`.

In both cases, the reaction may or may not remain part of the reactive program after the execution of the listener, depending on the reaction mode. If the mode is `ONCE`, the reaction is deregistered right after its execution. In particular, in the case of weak reactions we guarantee that the `ReactionListener` is executed only once, even when multiple matching tuples are returned from different hosts in the system. Instead, if the mode is `ONCEPERTUPLE` the reaction remains registered, but it must be prevented from firing multiple times for the same tuple. In the current implementation, this is accomplished by keeping a list of the tuple identifiers that have already been reacted to within the `RegisteredReaction` object itself. Each time a matching tuple is found, this data structure is queried and updated to determine if a listener should be executed. Our optimized implementation of `Reactor` separates newly written tuples from those that were in the tuple space prior to the execution of a given reaction, and greatly improves the performance of `ONCEPERTUPLE` by not selecting a tuple more than once from a single local tuple space. However, because tuples can migrate and weak reactions can be uninstalled and reinstalled as connectivity changes, it is possible for a tuple to be selected more than one for a reaction, making the list of tuple identifiers necessary.

Transient Sharing in a Federated Tuple Space: Reactions as a Run-time Building Block. One interesting aspect of our design is that the management of blocking operations over a transiently shared space, including a federated one, is performed by relying on reactions, which are exploited not only by the programmer through the API, but also internally to the run-time support.

When a blocking query operation is issued on a federated tuple space, two things may happen. In the case where a matching tuple is found immediately, the processing is very simple and coincides with the one taking place for non-blocking operations: the `LimeTSMgr` simply releases the calling agent and returns the matched tuple to it. Instead, if no matching tuple exists, further processing must detect when a matching tuple appears in the tuple space before being able to release the waiting agent. Essentially, the run-time support needs to *react* to the presence of a matching tuple—which is achieved precisely by using `LIME` reactions. When a blocking operation does not immediately find a tuple matching the given template, it creates and registers a `ONCE` reaction with the same template. When a matching tuple is eventually inserted in the tuple space, the tuple is passed to the listener through the `ReactionEvent` input parameter—as in all reactions. The listener is system-defined, and its body performs different actions depending on whether the operation originally requested was a `rd`, in which case it returns the matching tuple to the suspended agent, or an `in`, in which case the tuple is first removed from the host-level tuple space.

This processing takes place both in the case of a host-level and a federated tuple space, with the only difference being that in the first case a strong reaction is registered, while in the second case a weak reaction is used. This complicates things when an `in` must be processed. In fact, in the first case the use of a strong reaction guarantees that the tuple cannot be withdrawn in between the detection of its presence and the execution of the listener, since these two steps are executed atomically. This guarantee does not hold when a weak reaction is used, and hence a

subsequent `inp` must be used to withdraw the tuple. Probes are implemented by sending the operation request to the appropriate LIME server based on the location parameters. There, it is served by the appropriate `LimeTSMgr`, and the result is returned on the reverse path. In our case, if the probe returns a tuple the agent is released. However, the tuple might have been withdrawn by another agent, in which case the `inp` returns `null`, the reaction is re-registered, and the agent continues to wait.

Connectivity, Engagement, and Disengagement. Transient sharing of host-level tuple spaces into a federated one is dependent on the connectivity among the hosts in the system. A set of connected hosts forms a *group*. Independent, disjoint groups can co-exist, as long as they are not within communication range. When two or more groups move within range, they are merged into a single group containing all members. To provide this functionality, LIME relies on the group management protocol described in [29]. This protocol guarantees that, in response to a change in group membership detected at the network level, all members of the group are notified atomically. In principle, any group formation policy can be exploited, but the current implementation matches the above description.

In LIME, this notification triggers an update to the content of the `LimeSystem` tuple space, to reflect the new configuration of hosts. The other information concerning the configuration of the system, i.e., the agents residing at each host, and the names of the tuple spaces shared by each agent, are instead updated through the engagement and/or disengagement process, which is triggered by a group membership change and executed in the same atomic step. It should be noted that in the implementation, the `LimeSystem` is not actually implemented as a federated tuple space. Instead it is independently maintained at each host, and is kept consistent through updates during engagement and disengagement. Besides updating the information pertinent to the `LimeSystem`, engagement determines the migration of all the misplaced tuples whose destination became available during the group change. Moreover, for all the weak reactions involving newly connected hosts, engagement also triggers the registration of these reactions on the appropriate hosts. On the other hand, disengagement does not involve any transfer of information. Upon a group change weak reactions involving hosts that are now disconnected are simply deregistered, and the `LimeSystem` updated. No movement of tuples is required although, to preserve atomicity, the protocol we employ to perform the distributed transaction ensures that all hosts complete the disengagement process before regular operations are resumed.

Mobile Agents. LIME is a coordination framework that deals with mobility in a way that is independent from the nature of migration. As this is not one of its goals, it does not support directly agent mobility. Instead, as with tuple spaces, agent migration is decoupled from the rest of the system by an adaptation layer that simplifies the integration of a mobile agent system. The currently available implementation relies on an adaptor built for the `μCODE` mobile code toolkit, developed by one of the authors [22] and available as open source [20].

This adaptation layer allows a mobile agent to carry along one or more LIME tuple spaces, and automatically deals with their (dis)engagement. Upon a migration, the agent tuple spaces are all toggled to private, and hence disengaged. These tuple spaces are serialized as part of the agent state, and migrated to destination along with the agent, where they are deserialized and shared again before the agent code begins to execute. More details about the adaptation layer and how to integrate a mobile agent system with LIME are available in the LIME documentation, and on the LIME Web site [34].

Implementation Details. The `lime` package is roughly 5,000 non-commented source statements, resulting in an approximately 100 Kbyte `jar` file. The `LIGHTS` lightweight tuple space implementation and the adapter for integrating multiple tuple space engines adds an additional 20 Kbyte of `jar` file. When using mobile agents, the `μCODE` toolkit adds approximately 30 Kbyte in a `jar` file. Communication is completely handled at the socket level, requiring no support for RMI or other communication mechanisms. Thus far, LIME has been tested successfully on PCs running various versions of Windows and Linux, and exploiting both wired Ethernet as well as IEEE 802.11 wireless technology. Moreover, LIME runs successfully also on PDAs equipped with PersonalJava.

5 Developing Mobile Applications with LIME

Application development is the last phase of our research strategy, and the one where the abstractions inspired by formal modeling and embodied in the middleware are evaluated against the real needs of practitioners. In this section we present two applications that exploit the current implementation of LIME in a setting where physical

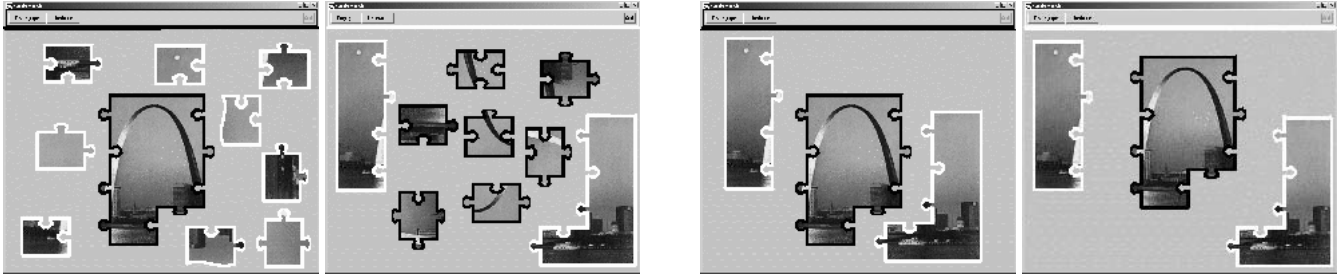


Figure 6: ROAMINGJIGSAW. The left two images show the puzzle trays of the black and white players while they are disconnected and able to assemble only their selected pieces. The right two images show the black and white puzzle trays after the players re-engage and see the assemblies that occurred during disconnection.

mobility of hosts is enabled. The first one involves the ability to perform collaborative tasks in the presence of disconnection, while the second one revolves around the ability to detect changes in the system configuration. In each case, we present the corresponding application scenarios and a report about the way LIME has been exploited during development. The lessons learned from these experiences and the results of our empirical evaluation of LIME are presented in the next section.

5.1 ROAMINGJIGSAW: Accessing Shared Data

Scenario. Our first application, ROAMINGJIGSAW, is a multi-player jigsaw assembly game. A group of players cooperate in a disconnected fashion on the solution of the jigsaw puzzle. They can construct assemblies independently (e.g., while disconnected), and share intermediate results or acquire pieces from each other when connected. Play begins with one player loading the puzzle pieces into a shared workspace that is visualized by the user as a *puzzle tray*. The workspace is shared among all connected users, therefore the puzzle trays of all users show the same set of puzzle pieces at this point.

Players can select pieces in the puzzle tray by clicking on them. The visual effect is that the piece outline is highlighted on all users' displays with the color of the selecting player. Selection has deeper consequences. In fact, although all the puzzle pieces are displayed on the tray, a player can make assemblies using only the pieces that she has selected, and that are currently displayed with her color. A player can select pieces or assemblies that are currently selected by another player, provided that the target player is connected.

Disconnection of a player does not have an immediate effect on the puzzle tray of the others. Nevertheless, pieces that have been selected by the departing player can no longer be selected by the others—and vice versa. Hence, the disconnected player can now construct assemblies by using only the pieces outlined with her color. Nevertheless, the pieces of all players remain visible. The assemblies made by each player during disconnection become visible to the others when connectivity among the players is restored. At this point, the view provided by the user interfaces is reconciled with the changes made during disconnection, and the selection of a piece belonging to a connected player is again possible. Figure 6 shows the appearance of the puzzle tray during disconnection and after reconnection.

From the description, it is evident that ROAMINGJIGSAW embodies a pattern of interaction where the shared workspace displayed by the user interface of each player provides an accurate image of the state of all connected players, but only a weakly consistent image of the global state of the system. For instance, a user's display contains only the last known information about each puzzle piece in the tray. If two pieces have been assembled by a disconnected player, this change is not visible to others. However, this still allows the players to work towards achieving the global goal, i.e., the solution of the puzzle, through incremental updates of their local state.

ROAMINGJIGSAW is a simple game that nonetheless exhibits the characteristics of a general class of applications in which data sharing is the key element. Hence, the design strategy we exploited in ROAMINGJIGSAW may be adapted easily to handle updates in the data being shared by real applications. One example could be provided by collaborative work applications involving mobile users, where our mechanism could be used to deal with changes in sections of a document, or with paper submissions and reviews to be evaluated by a program committee.

Design and Implementation. In our design of ROAMINGJIGSAW, we chose to represent pieces and assemblies as tuples, and the shared workspace as a tuple space. When a player selects a piece, the corresponding tuple is with-




Figure 7: REDROVER. The main console of REDROVER, the identifying icon of one of the team members, and one player (a ghost) out of range.

drawn and subsequently reinserted in the tuple space, with the field indicating the current “owner” automatically changed by LIME. Similarly, when a player builds an assembly out of several pieces, a new tuple is written containing information about the assembled pieces; the tuples associated with the latter are removed from the tuple space.

The critical issues in the design of ROAMINGJIGSAW are the detection of piece selection and assembly, the reconciliation of the puzzle tray taking place on reconnection, and the joining of a new player. Interestingly, all of these rely upon a single weak reaction of type `UbiquitousReaction` and mode `ONCEPERTUPLE`. Registration of the reaction is specified so that its template looks for any new tuple corresponding to a puzzle piece, while its listener takes care of updating the puzzle tray by using the information found in the tuple, thus correctly maintaining the weakly consistent view of the workspace. Since the reaction type sets its scope to the whole federated tuple space, the application receives updates about new pieces regardless of where and why they have been inserted, and hence notably without any need to be explicitly aware of the arrival and departure of players. Thus, the programming effort can be rightfully spent on handling data changes, rather than monitoring the system configuration.

Although the processing described thus far operates on the federated tuple space, fine-grained control over the location of tuples is critical in dealing with disconnections. To ensure that a player can access her selected pieces during a disconnection period, piece selection should actually transfer the corresponding tuple into the local tuple space of the player’s application. Moreover, according to what we discussed earlier, a player must be prevented from selecting a piece that is currently not present in the federated tuple space. For this reason, selection is performed by the application agent by issuing an `inp` operation on the tuple space of the player last known to have the piece. If the piece is returned, it is reinserted in the local tuple space of the new owner, thus leading to a successful selection. Otherwise, if no tuple is returned it means that the piece is unavailable for selection, and a message is displayed to the user.

5.2 REDROVER: Exploiting Context-Awareness

Scenario. Our second target application is a spatial game we refer to as REDROVER, in which individuals equipped with small mobile devices form teams and interact in a physical environment augmented with virtual elements. This forces the participants to rely to a great extent on information provided by the mobile units and not solely on what is visible to the naked eye.

REDROVER is the initial step in the development of a suite of games exploiting augmented reality. Our final aim is to empower each player with global positioning system (GPS) access, audio and video communication, range finding capabilities, and much more. Currently, the game is limited to seeking the physical flag of another team and gathering around the player who finds the flag. A snapshot of the graphical user interface is shown in Figure 7. The most dominant display element is a view of the playing field, indicating the current position of all players within range. In ROAMINGJIGSAW, disconnection was masked until a user tried to access a piece from a disconnected user. In REDROVER disconnection is instead made explicit, by displaying a “ghost” icon for a disconnected user, indicating her temporary unavailability.

Each player has an icon of herself that is available to her teammates, as shown in Figure 7. While this is currently a static image, the functionality can easily be extended to share recent images taken from a digital camera in order to share environmental information among teammates.

It is important to note that while location information should be disseminated to all players, this should not hold

for some of the data shared among the members of a team. For example, it is desirable to inform only one's own team members regarding the flag location, and not the opponents' team.

As with ROAMINGJIGSAW, REDROVER is a simple game with the potential to be extended to real world scenarios. Examples include the exploration of an unknown area by a group of people or robots. Our current efforts include the incorporation of a mapping mechanism, enabling users to define the elements of a region and share these results as they meet other users.

Design and Implementation. A key issue in REDROVER is to disseminate information about the physical location of a given user. In our implementation, this information is encoded as a tuple. The way tuples are filled with location information, and hence the back-end of the location support, is decoupled from the rest of the application through an appropriate API, and thus can easily accommodate various positioning alternatives. Currently, we support feeding of location data from a GPS device.

Instead of "pushing" the location information to other users, in our design we let each application agent "pull" it from the tuple space of the other users. At startup, all application agents register a weak ONCEPERTUPLE ubiquitous reaction, whose pattern matches any tuple containing location information. Each time a player moves, its location is updated by writing a tuple containing the corresponding information in the tuple space. This causes the reaction listener of all application agents (including the one corresponding to the moving user) to fire, and update the display according to the new location.

Detection of player disconnection relies on the LimeSystem tuple space and its support for reactions. In this tuple space, tuples corresponding to hosts and agents contain a field reporting the current connection status of the mobile unit. In REDROVER, all agents are registered for tuples in the LimeSystem representing departed hosts: the corresponding listener updates the display by substituting the icon of the player with the ghost icon. Similarly, when a player reconnects a complementary reaction fires that changes her icon back to normal. Interestingly, no further action is needed to retrieve the current location of the reconnected user, thanks to the ONCEPERTUPLE reaction associated to location tuples. In fact, if the player did not change its position while disconnected, its location tuple will still be the same and no reaction will occur. On the other hand, if the location tuple is different at engagement time, then the corresponding listener will be fired, updating the display.

To enforce a separation between team data and location information, we exploited LIME's ability to define tuple spaces with different names, and hence shared independently. Therefore, general information such as player location is written to a game-wide tuple space, while team-specific information like flags, or images is written to a tuple space accessible only by the members of that team. A player can either request to be notified when a given piece of information appears in the team tuple space, or can probe a specific player for a specific object. For instance, flag capture notification is implemented with a ONCE weak reaction over the team federated tuple space. Instead, the feature allowing to retrieve the player's icon is implemented through a rdp that goes directly to the tuple space of the corresponding remote player.

6 Discussion

In this section we elaborate on what we presented thus far. First, we report about lessons learned from the design and development of LIME. Then, we broaden the scope of our discussion and report about other projects that either build on or are inspired by LIME. Finally, we turn our attention to efforts that are not directly related with LIME, and survey related research projects that deal with tuple spaces and middleware for mobility.

6.1 Reflections and Lessons Learned

LIME is the result of a continuous interplay among the definition of the underlying formal model, the design and implementation of the middleware, and its evaluation on mobile applications. The development of a model for LIME, and its formalization, favored a better understanding of the abstractions provided by the middleware. In particular, by keeping the programming interface as close as possible to the operations defined in the formal model, we made it easy to communicate and reason about the functionality of the system and its use in applications. In an incidental way, this task also provided an evaluation of the applicability of Mobile UNITY to the specification of a middleware for mobility. The ability to think about abstractions in a setting unconstrained by implementation details favored a

style of investigation characterized by a more radical perspective, where the decisions driving the modeling and the definition of the main abstractions were mostly determined by the need for expressiveness and completeness.

Building applications on top of LIME made it possible for us to evaluate the usefulness of its programming abstractions and constructs. For example, our experiences confirmed the usefulness of weak reactions on the federated tuple space and the ONCEPTUPLE mode provide the programmer with a highly effective constructs that simplifies programming. The registration of a single reaction is sufficient to guarantee notification of relevant data as it appears throughout the federation, independent of changes in configuration. Interestingly, this power has a cost: the implementation of weak reactions is probably the most complicated portion of the current LIME software—this should be expected, since we are shifting a great deal of complexity away from the programmer and into the run-time support.

Another interesting by-product of these empirical evaluations is an understanding of the programming and architectural styles fostered by LIME and recurring in mobile applications. One distinction can be made between applications such as ROAMINGJIGSAW whose main requirement is to enable sharing of data *despite* mobility and those such as REDROVER where most of the computation is driven by reactions to changes in context and whose functionality *exists because of* mobility. In this and other application typologies, a recurring dilemma is between an application style that provides a weakly consistent view of the system in the presence of mobility, and one that provides a fully consistent view that takes into account departure and arrival of mobile units. Choosing one representation style or the other has non-trivial implications on the complexity of the overall design and development task, and on the primitives that must be used. If weak consistency is enough, the view can be built incrementally by exploiting the data notification mechanism provided by weak reactions, usually in the ONCEPTUPLE mode. If, instead, a fully consistent view is required, application specific machinery must be written in addition to using the LimeSystem tuple space to react (immediately) to changes in the system configuration. In our experience both styles are naturally accommodated by the abstraction of a transiently shared tuple space and use of the LimeSystem tuple space. Our “developers”, mostly graduate and undergraduate students, found it easy not only to *program* applications with LIME but, most importantly, to *think* about the application in terms of the metaphors characteristic of the underlying LIME model.

Despite the somewhat limited experience, analysis of LIME application developments revealed that the programming style induced by LIME is quite different from what we initially expected. This is especially true in the case of weak reactions and the LimeSystem tuple space, both of which prominently appeared in the ROAMINGJIGSAW and REDROVER application designs. Because LIME was envisioned to be a coordination framework founded on the idea of transiently shared tuple spaces accessible exclusively through Linda operations, reactive programming was not even part of the initial core. Similarly, the LimeSystem was initially thought of as an add-on to support very specific needs. Actually, we initially thought that the explicit context knowledge provided by the LimeSystem tuple space could be bypassed by the observation of changes in the data context. Experience with our applications, especially REDROVER, showed that this does not hold in general and that developers use extensively the LimeSystem tuple space.

6.2 LIME Inspired Projects

Although this paper presents the first comprehensive description of LIME, the published model and implementation have already been influential in multiple projects both inside and outside our research group.

Within our group, we have extended or exploited LIME for several purposes. First, we have begun to explore the issues of security in tuple space based ad hoc mobile environments [28] by allowing applications to protect selected tuple spaces and even individual tuples through the use of passwords. The same passwords were also used to encrypt communication among hosts when exchanging messages related to sharing specific tuples spaces. Second, we have used LIME as the foundation for a Jini-like service discovery mechanism [11]. This project, implemented as an application layer on top of LIME uses the tuple space for sharing service advertisements and performing pattern-based service discovery. This extends the client-server model of service discovery for the mobile ad hoc environment by coupling the services available for discovery with the services available in the network, and maintaining this connection even as connectivity changes. Finally, in another project the LIME tuple space is used to support code mobility by storing Java class bytecode [23]. The class loading mechanism is extended to resolve class names by searching the federated tuple space, instead of a well-known, centralized code repository. This mechanism enables the code on demand paradigm for code mobility in the mobile ad hoc environment, where connections to specific code servers are not always available.

Groups at other universities have presented alternatives to both the LIME implementation and the formalization.

At Purdue University, a group extracted the features of LIME necessary for mobile agents by removing host-level sharing, and created a model referred to as CoreLIME [6]. On top of this restricted model, they proposed some initial ideas for tuple space security. A group the University of Bologna proposed an alternative to the state-based formal specification presented here, providing a calculus-based specification [4]. This calculus presents several alternatives to the original LIME model, including reacting to tuple space operations instead of tuple space contents and blocking agents that generate tuples destined for disconnected agents rather than creating misplaced tuples.

As we conclude this section, it should be noted that the effort that went into developing LIME also contributed to the emergence of a more abstract and general coordination concept and methodology called *Global Virtual Data Structures* (GVDS) [25]. It is centered on the notion of constructing individual programs in terms of local actions whose effects can be interpreted at a global level. A LIME group, for instance, can be viewed as consisting of a global set of tuples and a set of agents that act on it in some constrained manner. The set has a structure that changes in accordance with a predefined set of policies and it is this very structure that governs the specific set of tuples accessible to an individual agent through its local interface at any given point in time. The analogy to the concepts of virtual memory and distributed shared memory are very strong and several other research projects have picked up the GVDS theme and instantiated it in their own unique ways. The XMIDDLE [14] system developed at University College of London, for instance, presents the user with a tree data structure based on XML data. When connectivity becomes available, trees belonging to different users can be composed, based on the node tags. Upon disconnection, operations on replicated data are still allowed, and their effect is reconciled when connectivity is restored. Also PEERWARE [9], a project at Politecnico di Milano, exploits a tree data structure, albeit in a rather different way. In PEERWARE, each host is associated with a tree of document containers. When connectivity is available, the trees are shared among hosts, meaning that the document pool available for searching under a given tree node includes the union of the documents at that node on all connected hosts.

6.3 Related Projects

The last several years have seen a revitalization of Linda for distributed computing applications, including mobile environments. From the industrial perspective, both Sun and IBM have developed tuple space implementations for client-server coordination, i.e., JavaSpaces [13] and TSpaces [12], respectively. These systems present a centralized tuple space, accessible through remote operations by multiple processes. Their client-server approach is significantly different from that of LIME, as we target mobile ad hoc applications with implicit access to the remote data of other hosts, and support mobile agents.

Distributed Linda implementations have been studied extensively for fault tolerance [36, 2] and data availability [27]. The main disadvantage with these approaches is their need for high degrees of connectivity among the hosts of the distributed portions of the tuple spaces, a property inherently not present in the mobile environment.

One of the first applications of Linda to mobility came in the Limbo platform [3, 35], a system that builds the notion of quality of service aware tuple spaces that reside on mobile hosts. The quality of service information itself is stored in the tuple spaces and can be made accessible to agents on remote hosts. While Limbo has a notion of distributed tuple spaces that span multiple hosts, there are no mobile agents carrying tuple spaces when they migrate, no concept of reaction, and the mechanism for relocation of tuples is unclear. Interestingly, the Limbo *universal tuple spaces*, which serves as a registry for all tuple spaces is similar to the LimeSystem tuple space of LIME. However, instead of describing the *current* system context, the universal tuple space remembers all tuple spaces the host has ever encountered without regard for current reachability.

Two other models, TuCSoN [19] and MARS [5], exploit tuple space coordination for mobile agents, creating *programmable tuple spaces*. When an agent poses a query to the tuple space, the registered reaction that matches the operation fires, and an action is atomically performed. While in LIME reactions form a core concept for the application programmer, MARS and TuCSoN reactions are designed to be implemented by the system support designer to provide an intermediate access between the form of the query (which can vary among agents) and the data (which remains constant within a host, but is adapted when a query arrives). Another feature of MARS and TuCSoN is the option to fully qualify a tuple space name, identifying the specific host where the tuple space resides. This enables remote operation on tuple spaces, but connectivity must be available and the agent must be explicit about interaction, as opposed to the LIME model that operates over the current context transparently. Further, in MARS and TuCSoN, mobile agents only have access to the tuple spaces fixed at the hosts, they do not carry tuples as they migrate, and there is no coordination or data exchange among tuple spaces.

The KLAIM [18] model supports a programming paradigm where code migrates during execution, using tuple

spaces to provide the medium for interaction among processes. Tuple spaces have locality, but unlike in LIME, these tuple spaces are not permanently associated to a process. Instead, KLAIM processes located at a given locality implicitly interact through the co-located tuple space. There is no transient sharing among tuple spaces, but a process can explicitly interact with any tuple space by identifying its locality, and a process can migrate to a new locality to interact locally. While LIME leaves the details of process migration outside the model, KLAIM includes in the formal specification the details of process migration, making it an integral part of the model.

As alluded to in the informal description of LIME provided in Section 2, the notion of reaction put forth in LIME is profoundly different from similar event notification mechanisms such as those provided by TuCSoN, TSpaces, and Javaspaces. In these systems, the events respond to *operations* issued by processes on the tuple spaces (e.g., out, rd, in, etc.). In LIME, however, reactions fire based on the *state* of the tuple space itself. Further, LIME reactions execute as a single atomic step, and cannot be interrupted by other operations. This makes it straightforward for a single LIME reaction to probe for a tuple, react if it is found, and register a reaction if it is not. This same operation in the other systems requires a transaction. Finally, the atomicity of strong reactions increases the power of LIME reactions. For example, with a strong, local reaction, the execution of the listener is guaranteed to fire in the same state in which the matching tuple was found. No such guarantee can be given with an event model where the events are asynchronously delivered, nonetheless, we support this second approach through weak reactions.

7 Conclusions

LIME is a middleware specifically designed to support logical mobility of agents and physical mobility of hosts in both wired and wireless settings. Within this general context, its distinctive feature is the reliance on coordination to simplify the development of mobile applications. While building on the decoupling advantages of the original Linda model, LIME breaks new ground by extending coordination technology to mobile systems, including the ad hoc wireless setting. Transparent management of tuple space sharing, contingent of connectivity, offers an effective context awareness mechanism while reactions provide an effective and uniform vehicle for responding to context changes regardless of their nature or trigger. The net result is a simple model with precise semantics and applicability in a wide range of settings, from mobile agent systems operating over wired networks, at one extreme, to ad hoc networks lacking any infrastructure support, at the other. While a full formal validation of LIME's impact on software development productivity is still to be performed, our experience to date with the development of a reasonable set of applications in wireless settings appears to validate LIME's potential for rapid development of mobile applications. Finally, it should be noted that LIME, in addition to demonstrating the practical use of coordination technology in mobile computing, opens a new area of research involving the application of state-based coordination models and middleware to context-aware computing.

Availability. LIME continues to be developed as an open source project, available under GNU's LGPL license. Source code and development notes are available at lime.sourceforge.net.

Acknowledgments

The authors wish to thank Pietro Braione for his insightful comments on early drafts of this paper. This research was supported in part by the National Science Foundation under grant No. CCR-9970939. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the research sponsors.

References

- [1] R.J.R. Back and K. Sere. Stepwise Refinement of Parallel Algorithms. *Science of Computer Programming*, 13(2-3):133–180, May 1990.
- [2] D.E. Bakken and R. Schlichting. Supporting fault-tolerant parallel programming in Linda. *IEEE Transactions on Parallel and Distributed Systems*, 1994.
- [3] G. Blair, N. Davies, A. Friday, and S. Wade. Quality of Service Support in a Mobile Environment: An Approach Based on Tuple Spaces. In *Proc. of the 5th IFIP Int. Wkshp. on Quality of Service (IWQoS'97)*, May 1997.

- [4] N. Busi and G. Zavattaro. Some thoughts on transiently shared dataspace. In *Proc. of the Workshop on Software Engineering and Mobility, co-located with the 23rd Int. Conf. on Software Engineering ICSE*, 2001.
- [5] G. Cabri, L. Leonardi, and F. Zambonelli. Mars: A programmable coordination architecture for mobile agents. *IEEE Internet Computing*, 2000.
- [6] B. Carbutar, M.T. Valente, and J. Vitek. LIME revisited: Reverse engineering an agent communication model. In *International Conference on Mobile Agents*, Atlanta, GA, USA, December 2001.
- [7] N. Carriero, D. Gelernter, and L. Zuck. Bauhaus-Linda. In *Workshop on Languages and Models for Coordination, European Conference on Object Oriented Programming*, 1994.
- [8] K.M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.
- [9] G. Cugola and G.P. Picco. PEERWARE: Core middleware support for peer-to-peer and mobile systems. Technical report, Politecnico di Milano, Italy, 2001. Available at www.elet.polimi.it/upload/picco.
- [10] D. Gelernter. Generative Communication in Linda. *ACM Computing Surveys*, 7(1):80–112, Jan. 1985.
- [11] R. Handorean and G.-C. Roman. Service provision in ad hoc networks. In F. Arbab and C. Talcott, editors, *Proceedings of the 5th International Conference on Coordination Models and Languages*, LNCS 2315, pages 207–219. Springer, 2002.
- [12] IBM. TSpaces Web page. <http://www.almaden.ibm.com/cs/TSpaces>.
- [13] JavaSpaces. The JavaSpaces Specification web page. <http://www.sun.com/jini/specs/js-spec.html>.
- [14] C. Mascolo, L. Capra, S. Zachariadis, and W. Emmerich. XMIDDLE: A data-sharing middleware for mobile computing. *Kluwer Personal and Wireless Communications Journal*, 21(1), April 2002.
- [15] P.J. McCann and G.-C. Roman. Compositional Programming Abstractions for Mobile Computing. *IEEE Trans. on Software Engineering*, 24(2):97–110, 1998.
- [16] P.J. McCann and G.-C. Roman. Modeling Mobile IP in Mobile UNITY. *ACM Transactions on Software Engineering and Methodology*, 8(2):115–146, April 1999.
- [17] A.L. Murphy, G.P. Picco, and G.-C. Roman. LIME: A Middleware for Physical and Logical Mobility. In F. Golshani, P. Dasgupta, and W. Zhao, editors, *Proc. of the 21st Int. Conf. on Distributed Computing Systems (ICDCS-21)*, pages 524–533, May 2001.
- [18] R. De Nicola, G. Ferrari, and R. Pugliese. KLAIM: A Kernel Language for Agents Interaction and Mobility. *IEEE Trans. on Software Engineering*, 24(5):315–330, 1998.
- [19] A. Omicini and F. Zambonelli. Tuple Centres for the Coordination of Internet Agents. In *Proc. of the 1999 ACM Symp. on Applied Computing (SAC'00)*, February 1999.
- [20] G.P. Picco. μ CODE Web page. mucode.sourceforge.net.
- [21] G.P. Picco. LIGHTS Web page. lights.sourceforge.net.
- [22] G.P. Picco. μ CODE: A Lightweight and Flexible Mobile Code Toolkit. In *Proc. of the 2nd Int. Workshop on Mobile Agents*, LNCS 1477. Springer, 1998.
- [23] G.P. Picco and M.L. Buschini. Exploiting transiently shared tuple spaces for location transparent code mobility. In F. Arbab and C. Talcott, editors, *Proc. of the 5th Int. Conf. on Coordination Models and Languages*, LNCS 2315, pages 258–273. Springer, 2002.
- [24] G.P. Picco, A.L. Murphy, and G.-C. Roman. LIME: Linda Meets Mobility. In D. Garlan, editor, *Proc. of the 21st Int. Conf. on Software Engineering*, pages 368–377, May 1999.
- [25] G.P. Picco, A.L. Murphy, and G.-C. Roman. On global virtual data structures. In D. Marinescu and C. Lee, editors, *Process Coordination and Ubiquitous Computing*, pages 11–29. CRC Press, 2002.
- [26] G.P. Picco, G.-C. Roman, and P.J. McCann. Reasoning about Code Mobility with Mobile UNITY. *ACM Trans. on Software Engineering and Methodology (TOSEM)*, 10(3):338–395, July 2001.
- [27] J. Pinakis. *Using Linda as the Basis of an Operating System Microkernel*. PhD thesis, University of Western Australia, Australia, August 1993.
- [28] G.-C. Roman and R. Handorean. Secure Sharing of Tuple Spaces in Ad Hoc Settings. Technical Report WUCS-02-31, Dept. of Computer Science and Engineering, Washington Univ. in St. Louis, MO, USA, 2003.
- [29] G.-C. Roman, Q. Huang, and A. Hazemi. Consistent group membership in ad hoc networks. In *Proceedings of the 23rd Int. Conf. on Software Engineering*, pages 381–388, Toronto, Canada, May 2001.
- [30] G.-C. Roman, P.J. McCann, and J.Y. Plun. Mobile UNITY: Reasoning and specification in mobile computing. *ACM Transactions on Software Engineering and Methodology*, 6(3):250–282, 1997.

- [31] G.-C. Roman, A.L. Murphy, and G.P. Picco. Coordination and Mobility. In A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors, *Coordination of Internet Agents: Models, Technologies, and Applications*, pages 254–273. Springer, 2000.
- [32] D.S. Rosenblum and A.L. Wolf. A Design Framework for Internet-Scale Event Observation and Notification. In *Proc. of the 6th European Software Engineering Conf. held jointly with the 5th ACM SIGSOFT Symp. on the Foundations of Software Engineering (ESEC/FSE97)*, number 1301 in LNCS, Zurich (Switzerland), September 1997. Springer.
- [33] A. Rowstron. WCL: A coordination language for geographically distributed agents. *World Wide Web Journal*, 1(3):167–179, 1998.
- [34] Lime Team. LIME Web page. lime.sourceforge.net.
- [35] S.P. Wade. *An Investigation into the use of the Tuple Space Paradigm in Mobile Computing Environments*. PhD thesis, Lancaster University, England, September 1999.
- [36] A. Xu and B. Liskov. A design for a fault-tolerant, distributed implementation of Linda. In *Digest of Papers of the 19th Int. Symp. on Fault-Tolerant Computing*, pages 199–206, June 1989.