Limit the Space of Agent Interaction to Friends

Matteo Ceriotti
FBK-IRST & University of Trento, Italy, ceriotti@itc.it

ABSTRACT

LIME is a tuple space supporting the coordination of agents. Each agent owns a part of the globally shared knowledge, which is represented by a set of tuples, for which the agent is responsible. An agent can share its own tuples with the other agents in the range of communication. A new component of the middleware allows the definition of application-specific policies that restrict the scope of the sharing of the context of an agent. Any policy can be defined and it is applied when an operation is issued on the tuple space; the result is guaranteed to reside in the local tuple space of an agent which agrees on the sharing rules defined by the one that submits the request (the agent that replies can also reject the request on the basis of its own policy). Different policies are allowed to coexist and cooperate in the same environment.

RESPECT is a language that defines the behavior of a tuple centre, aimed at coordinating the interactions among the agents. The tuple centre is an artifact that can be accessed by any agent; the coordination laws constrain the interactions among the agents but they do not bind to the desires of the agent itself. However, each agent has its own rules and preferences in selecting the "friends" with which collaborates and also the coordination medium can take advantage from the possibility to restrict the access to the context only to a group of agents which agree on some specific behavior constraints.

In this work, I present a possible mapping between the LIME and the RESPECT model; the evidence is put on the mechanisms that allow the restriction of the space of interaction among the agents from both the point of view of the agents and of the coordination medium.

1. INTRODUCTION

An agent is an autonomous computational entity that encapsulates control and the rules that govern it. It differs from both an artifact, which is aimed to be used by the agents, and the environment, which is the surrounding of both the agents and the artifacts. These three elements together with their interactions are defined as a multi-agent system. In this context, the agents are usually not alone and one of the main issues of such a system is to allow the coordination and the communication among the agents. Whereas the result of the interaction with an artifact is predictable, the communication and the coordination among autonomous agents lead in general to unpredictable outcomes due to the nature of the agent itself.

One of the main characteristics of the agents is their sociality. This property comes directly from their autonomy, which can be completely exploited only inside a society of agents. In natural societies, the individuals form groups based on the goals that the group allows the member to reach. These goals are usually different inside each single group and the collaboration among the agents in one group is characterized by some goals that can be reached with the contribute of each single member. To improve the feasibility of a purpose, the individuals group based on some rules that are defined both explicitly by each agent and implicitly by the environment. In this way, the communication is restricted among the agents that are inside the scope defined by these rules. We can consider the parallel with the relationship among friends where both the rules defined by the environment and the one defined by each member contribute to restrict the access to the ones who respect the laws of the group.

A great effort has been spent in defining platforms to support the communication and the coordination among the agents; one model that fits the requirements introduced by the A&A (agents and artifacts) infrastructure is the tuple space. This coordination medium offers nice properties that make the interactions among agents feasible and suitable. Several extensions of the model have been proposed with specific motivations and goals; among them LIME, which aims to face logical and physical mobility of the agents, and RESPECT which provides a language to define the behavior of the coordination medium.

Recently, the introduction of a new feature inside the LIME middleware has been discussed. Its main purpose is to define the scope of the sharing of the information; by defining that, the new component allows the grouping of agents based on some defined policies. In this work, I discuss the possibility to introduce this extension in the RESPECT model.

The work is structured basically in the following way: the Section 2 introduces the tuple space and the tuple centre models and provides a discussion about the possible matches between the two, in particular presenting LIME and RESPECT, and their analogies and dissimilarities; the Section 3 presents the extension which allows the restriction of the scope of the interaction and applies such feature to both the models; the last Section 4 provides the conclusion of the work with some possible future directions.

2. LIME AND RESPECT: THE COORDINATION BASED ON THE TUPLES

The communication among parallel processes based on tuple spaces was introduced the first time in Linda [4]. The concurrent programming model described there is called generative communication and it is based on the concept of tuple spaces, containers of ordered sequences of values called tuples. Its main characteristic is the fully distribution in both time and space; the producer and the consumer of data can be unknown to each other and they are not required to be active in the same time to the communication to be exploited.

The model introduced in Linda is essentially composed by three primitives: **out** (to insert a new tuple in the tuple space), **in** (to retrieve one tuple from the tuple space) and **rd** (to read one tuple without removing it from the tuple space). The data available in the tuple space are accessed by matching tuples against a given template. Given that the basic primitives are blocking (a result is returned only when a matching tuple is found), some extensions (**inp** and **rdp**) were introduced to allow results describing the absence of matching tuples.

In the field of the distributed systems, several middlewares supporting the tuple space model were defined, each one with its own distinguishing features. Among them, LIME (Linda In a Mobile Environment) [6] aims to face the mobility both logical and physical of agents and hosts by adapting the model to that scenario. Each host is responsible for a part of the global tuple space; when several hosts with equally named tuple spaces are in the range of communication they transiently share their tuples with the others. The introduction of reactive primitives allows the agents to react, not only to the insertion of new tuples, but also to the changes in the context due to the mobility of the nodes.

A new agent joins the shared context through the atomic operation of engagement. This operation makes the tuples of the joining agent available to the other agents and vice versa. Before unifying the two contexts, the one of the joining agent and the one already shared among the other actors, the operations already issued on the tuple space are evaluated so that if a matching tuples is found as a consequence of the joining of a new context the result is returned before any other operation is issued. In this way, the LIME model provides reactivity to the changes due to the mobility.

One of the domains in which the tuple space model offers great support is the field of multi agent systems. In general it can be defined as a computational system made of agents and artifacts; the former are autonomous computational entities which encapsulate control, the latter are computational entities aimed at the use by agents; these two elements are then immersed in the environment which is influenced both by the actions issued by the agents and by the effects of the utilization of the artifacts. In this scenario, the tuple space model meets the need for a coordination medium connecting the agents together. However, the kind of coordination fostered by the tuple space model lacks the possibility to define the behavior of the coordination media itself, in response to the interaction of the agents with the tuple space.

This lack motivates a new coordination medium, the tuple centre, which is an extension of the basic tuple space model that supports the definition of the coordination medium behavior. ReSpecT [5] is a language through which such behavior can be described. Its main features are the possibility to define computations inside the tuple centre through reactive primitives, called reactions, and the possibility to bind them not only to the availability of data (as the operations available in the tuple space model) but also to the interactions of an agent with the tuple centre, i.e. the operations issued. In this way, the tuple centre results in a composition of two parts: the tuple space, with ordinary tuples through which the communication is exploited, and the specification space, in which the coordination laws are described by the corresponding tuples.

The model of LIME and the one of ReSpecT are different, in particularly in the reactivity of the tuple space medium. In fact, while the coordination medium in the case of LIME reacts only to the changes in the topology of the nodes and the operations submitted by the agents entirely define the coordination, the tuple centre influences with its own behavior the interactions among the agents. LIME however faces the mobility of the agents by giving to each of them the responsibility for a part of the global data; in this way the system is aware of the components which are collaborating on the same tuple space. This is particularly interesting for the decentralization of the coordination that relies on the behavior and the availability of each single node.

In both the systems, the coordination happens among all the available agents, which are the ones in the range of the communication operating on the same tuple space in the case of LIME, or the ones that are interacting with the same tuple centre in the case of ReSpecT. In general, neither the agents nor the tuple centre can prevent an agent from accessing the data inside the shared context.

3. RESTRICT THE SCOPE OF THE INTERACTION

A new component is going to be introduced in the LIME middleware with the purpose of restricting the scope of the engagement operation. In heterogeneous environments, in fact, not all the agents can offer the same guarantees to the others; for example, an agent can move too fast to wait for all the installed operations to be activated against its own tuples. To obtain that, each single agent defines its own engagement policy that describes its own behavior and the agents with which it wants to share its part of the tuple space. In this way when a new agent is available or when it changes one of its attributes that describe its behavior, the engagement policies of the other available agents are evaluated against the one of that agent; if the engagement policies concord with one another, the engagement is activated and the agent which requested the engagement can operate on the information available in the context of the agents that have accepted the engagement.The engagement operation is asymmetric which means that each agent has its own particular view of the system, but the policy can require the engagement to be issued by both the involved entities so that a symmetric relationship is established.

This extension offers a more fine-grained control over the interaction, which is completely decentralized and based on the point of view of each single agent and on its behavioral description. This extension could lead to an anarchic system where any member defines its own rules; in reality it is prevented from happening because the communication and cooperation among agents is allowed only among agents which agree on the same rules and, as it can be easily seen, the more are the agents that agree on a policy, the more are the chances to communicate and retrieve information.

The model of the tuple centre provides one single coordination medium in which the tuples reside, shared among all the agents; the agents are not responsible for any tuple and the coordination laws are defined inside the tuple centre, outside the agent. No agreement is required to access the information and the agents are unaware of the members that are collaborating on the same medium. In general these properties are accepted in order to make the system open and to completely decouple the users of the tuple space which can cooperate without any explicit knowledge of the respective identities. However in several scenario, both the agents and the coordinating entity can take advantages from restricting the access to a given set of actors with given properties or requirements.

An example can be provided to better understand what the model gains by the definition of engagement policies. This example is based on the problem of the drinking philosophers [1] which is a generalization of the well know dining philosophers problem [2,3]. Each philosopher has its own set of bottles from which it wishes to drink and it may have different sets of bottles describing its own preferred "cocktails". A philosopher can drink only when it holds all the bottles necessary to make a cocktail (all the bottles in one set). For the purposes of this work, the main interesting points of this problem are the wishes of each philosopher that are in general different from the ones of the others and the fact that the philosopher can successfully drink only if the specific bottles necessary to satisfy its own taste are available. We can extend the problem by considering the bottle as a consumable resource; in that case, the philosophers cannot continue to drink without providing new bottles and the resources could not be enough for all the philosophers.

In the drinking philosophers problem the coordination is defined by common rules but are based on the flavors of the agents. While the coordination laws are still properties of the coordination medium, there are evaluations that both the agents and the tuple centre are required to take before the interaction can actually take place. In fact, an agent has no interest in seating at a table where there are not the bottles that can satisfy its own desire; in a similar way, a seat should not be offered to an agent for which enough resources cannot be guaranteed. Furthermore, during the evaluation of the access request, information can be given in such a way that the agent can take properly actions to make the access possible (i.e. bring one new bottle or change the wishes).

The extension of the ReSpecT model to support this feature is quite straightforward. Both the agent and the tuple centre offer a public description that can be used to evaluate the acceptance or the rejection of the access to the coordination medium. The agent, when provided with the description of the tuple centre, can decide if it is worth to use it or not. In a similar way, the tuple centre can reject the request of connection made by an agent; if the rejection rules can be defined a priori, it can be included in the definition of the coordinator behavior, otherwise the tuple centre can ask to the current members to reach an agreement on the acceptance of the joining agent. In this last case the coordination medium is the one that allows not only the coordination among the agents interacting with the medium itself but it also coordinates the agents in deciding the other agents with which the interaction actually is allowed; furthermore, being the coordinator it can also provide rules such that fairness is guaranteed from the agents already members of the group in responding to the engagement requests.

4. CONCLUSION

In this work, I provide a discussion about two coordination media for interactions among agents, the tuple space and the tuple centre. Two different implementations of the model are also described, LIME and ReSpecT, putting in evidence the differences between them. Indeed, I present the extension of the model exploited by ReSpecT with a feature already introduced in the LIME model; the purpose of such an extension is to provide the possibility to restrict the scope of the agent interaction with the definition of special engagement policies.

Some possible future directions can be highlighted here. It is interesting to analyze possible changes to ReSpecT in order to provide a completely decentralized and individual interaction model, as it is defined in the version of LIME augmented with the support to the engagement policies. As we described, the LIME and ReSpecT models are different; some of the main differences are the ownership of the data in the LIME model and the reactivity to the network topology; it is interesting also to think about the way the ReSpecT model can be extended to support such features.

5. REFERENCES

[1] K. M. Chandy and J. Misra. The drinking philosophers problem. *ACM Trans. Program. Lang. Syst.*, 6(4):632–646, 1984.

[2] E. W. Dijkstra. Hierarchical ordering of sequential processes. In *Operating Systems Techniques*, pages 72–93. 1971.

[3] E. W. Dijkstra. Two starvation-free solutions of a general exclusion problem. Circulated privately, 1977.

[4] D. Gelernter. Generative communication in linda. *ACM Trans. Program. Lang. Syst.*, 7(1):80–112, 1985.

[5] A. Omicini and E. Denti. Formal ReSpecT. *Electronic Notes in Theoretical Computer Science*, 48:179–196, June 2001. Declarative Programming – Selected Papers from AGP 2000, La Habana, Cuba, 4–6 Dec. 2000.

[6] G. P. Picco, A. L. Murphy, and G.-C. Roman. Lime: Linda meets mobility. In *ICSE '99: Proceedings of the 21st international conference on Software engineering*, pages 368–377, Los Alamitos, CA, USA, 1999. IEEE Computer Society Press.