Integrating a MAS and a Pandemonium: the Open-Source Framework AKIRA

Giovanni Pezzulo
ISTC-CNR
Viale Marx, 15
00137 Roma, Italy
+39 06 8609 0208
g.pezzulo@istc.cnr.it

Gianguglielmo Calvi
ISTC-CNR
Viale Marx, 15
00137 Roma, Italy
+39 06 8609 0208
gianguglielmo.calvi@noze.it

Rino Falcone
ISTC-CNR
Viale Marx, 15
00137 Roma, Italy
+39 06 8609 0211
r.falcone@istc.cnr.it

ABSTRACT
The open-source framework AKIRA integrates Multiagent and Pandemonium elements. We describe the main components of the framework, showing that the hybrid nature of the Agents, having symbolic and connectionist features, permits to model many functionalities such as implicit communication and coordination.

Categories and Subject Descriptors
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

General Terms
Algorithms, Design.

Keywords
Open-Source, Multiagent Systems, Pandemonium

1. INTRODUCTION
AKIRA [7] is an open-source, run-time C++ multithreading environment for building and executing Agents. We introduce its main peculiarities with respect to existing agent technology, such as the hybridization with the Pandemonium [1] model.

2. AN HYBRID MAS-PANDEMONIUM
AKIRA implements both a MAS and a Pandemonium [1]. The pandemonium was originally a bottom-up feature recognition model (e.g. for letter recognition). In its simplest form it consists in a set of Daemons, usually divided into layers; each Daemon is responsible for testing a condition of its environment; if it succeeds, it "shrieks" –and its shriek is used by Daemons of higher levels as input data e.g. for calling one of its (more complex) functions. The environment of a Daemon can either be the "real environment" (like sensors) or the activity of the Daemons at the lower level. In our implementation there are no explicit "layers": Daemons of whichever complexity can become "active" (the metaphor of "going from the stands into the playing field" is used there) and thus influence and spread activation to each other, providing that they have a "link". An appropriate taxonomy of the "links" makes it possible to model both top-down and bottom-up dynamics. In AKIRA the kernel represents the whole Pandemonium and the Agents represent the Daemons. The Pandemonium is responsible for many system operations such as initializing the system; launching, monitoring and killing the Daemons; garbage collecting; serving as "agent name server", etc. The Daemons are active objects: each one has its own thread of execution. Daemons are roughly the equivalent of Agents in MAS, but they include some connectionist features, too. Agents share messages through a Blackboard and two specialized shared structures, providing efficient peer-to-peer messaging and data sharing. We included a number of programming facilities and libraries, including BDI [8], fuzzy logic [3], neural networks [5]. The hybridization with the Pandemonium enhances the MAS design space, normally including only the Agent and System levels. At the contrary, AKIRA allows designers to exploit the emergent behavior of "groups of Agents" (Coalitions) that can be seen as super-agents, subsuming the behavior of many Daemons.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

2.1 Symbolic and Connectionist Elements
AKIRA Agents are Hybrid: they have a "symbolic" body (that can e.g. process rules and exchange symbolic messages); but they also have connectionist elements. AKIRA can thus be described at two levels of detail. At the level of connectionist elements, Daemons are nodes in a networks; they have an activation level and are linked each other via an Energetic Network that is a carrier of energy. At the level of functionalities, Daemons have a symbolic body that carries on a specific operation. The semantic of the two levels is interlaced: more energy at the connectionist level corresponds to more computational resources at the functional level: more computational resources, i.e. an higher priority of its thread and thus more time for its operations. This feature is called Energetic Metaphor in [2]. *Fig. 1* shows the two levels. Moreover, there exists a pool of shared resources, the Energy Pool, that gives an upper bound to the total resources (energy) of the system, making it conservative. Access to the resources is concurrent (serialized by Agents activation): if some energy is tapped by an Agent, it will be not available to the other Agents. Agents can spread their energy via the Energetic Network. Agents also release energy to the Energy Pool when they perform their symbolic operations, each having a "cost". This energetic model is called *AKIRA Energetic Model*, AEM [7]. The connectionist features allow programmers to design complex

"social" dynamics using Agents and Coalitions as actors. AKIRA extends the MAS perspective on distribution of the tasks and teams formation: it is not limited to explicit communication and coordination, but exploits connectionist dynamics such as auto-organization, implicit cooperation and concurrence, emergence of global dynamics from local rules. For example, differently from many role-based and symbolic team aggregation policies, Agents can form temporary or permanent teams, the Coalitions, even via energetic dynamics, i.e. without explicit messaging.

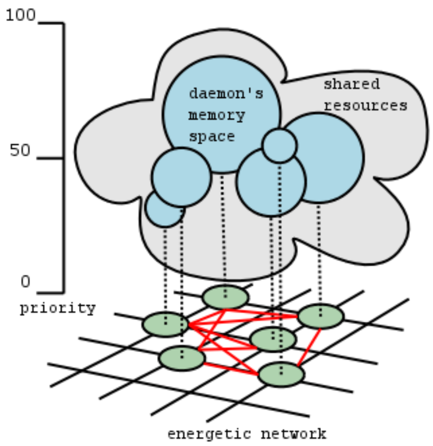

Figure 1. The two aspects of AKIRA

3. SOCIALITY

Agents can be involved in "social" interactions in many ways: a) by exchanging peer-to-peer messages via a Blackboard, a shared data structure where the messages are asynchronous and concurrently written and read; b) by sharing variables or objects through two specialized structures; c) by spreading activation, via the Energetic Network; d) by forming Coalitions. The Blackboard and the global factories are optimized "communication channels" for sharing explicit information in KQML-like format. The Blackboard can not only be used as a message dispatcher: it can also implement a "common workspace" where the Daemons notify their current activity and activation (even without an explicit receiver). This Pandemonium feature can be exploited for implicit communication: when they successfully perform an operation, Daemons notify (by shrieking) their success (and their level of activation) to the Blackboard. This notification can be interpreted as a "message" by other Daemons, that can even interpret or use this "message" in different ways, in their own terms, even without a shared semantic or ontology. This feature is used in Pandemonium models for bottom-up tasks, such as recognizing complex patterns by integrating a set of features, or collecting data to be used for a reasoning; it models the fact that the perceptual field of some Daemons is not the environment but the activity of other Daemons. This form of communication can be extended for more complex tasks, even top-down, especially if coupled with energetic dynamics as in AKIRA. As an example of a top-down dynamic, a Daemon can decide to spread some activation to other Daemons if it knows that it can exploit the results of their activity (if they succeed); this approach is used as the basis of the Behavior Networks [4], an action selection mechanism. Moreover, there can be "implicit" coordination dynamics: e.g. a Daemon can decide to collaborate with another one, e.g. by satisfying a condition that improves or makes it possible its work; or by sending it the extra energy it needs. These considerations lead to teamwork dynamics.

3.1 Team Work, Coordination and Coalitions

Many existing systems support the concept of team work: a task or a goal is assigned to a group of agents. Teams are widely used, e.g. for resolving a task together [6]; share plans or intentions [10]; exploring the issues of autonomy, norms, institutions, trust and control [9]. Team management techniques can vary (roles, hierarchies, distributed plans, specification of concurrency constraints), but in general they rely upon an explicit role/task assignment and use explicit messaging; often global state information and synchronized time is needed. In the connectionist literature a different concept of coordination exists: "implicit" coordination, including emergent behavior, self-organization and stigmergy. Recently this approach has been applied even to MAS operations, focusing on "implicit" forms of coordination, e.g. in particular stigmergic and mediated by the environment or coordination artifacts. In analogy with "implicit communication", a Pandemonium endorses "implicit coordination" in a very natural way, both for tasks requiring an explicit hierarchy of Agents and for "mediated" coordination. AKIRA Agents can form higher level assemblies called Coalitions, that are temporary sets of agents having energetic and symbolic exchange. Coalitions are not centrally controlled, nor are stable, but dynamically emerge (and collapse) depending on how much exchanges the single agents have. Agents can thus join or leave Coalitions, or move from one another. Coalitions afford an implicit form of collaboration between the Agents, by realizing tasks that are impossible to perform for each single one. However, their interaction is not predefined, but results from an emergent, context-sensitive process. Of course, an accurate design is necessary for managing parallelism and a lack of centralized timing and schedule (and for avoiding that Agents hinder each other rather than collaborate) [4]. But the advantage is that, differently from hierarchical structures, Coalition activity has not to be pre-planned[1].

4. REFERENCES

[1] Jackson J. V., Idea for a Mind. *Siggart Newsettler*, 181:23-26, 1987

[2] Kokinov B. N., The context-sensitive cognitive architecture DUAL, in *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum Associates, (1994).

[3] Kosko, B. Fuzzy cognitive maps. *International Journal of Man-Machine Studies*. 1986 (24)

[4] Maes P., Situated Agents Can Have Goals. *Robotics and Autonomous Systems*, 6 (1990).

[5] McClelland, J. L. & Rumelhart, D. E. (1988). *Explorations in Paralell Distributed Processing: A Handbook of Modles, Programs and Exercises*. MIT Press, Cambridge, MA.

[6] Minsky M. *The Society of Mind*. Simon and Schuster, N. Y. 1986

[7] Pezzulo, G. Calvi G. (2004). AKIRA: a Framework for MABS. *Proceedings of MAMABS 2004*

[8] Rao A., Georgeff M., BDI Agents from Theory to Practice, *Tech. Note 56, AAII*,1995.

[9] Castelfranchi, C. & Falcone, R. (in press). Social Trust Theory. In Pitt, J. *The open Agent Society*. John Wiley & Sons.

[10] Woolridge M. and Jennings N.R.. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115--152, 1995

[1] Work supported by the EU project **MindRACES**, FP6-511931.