

Distributed Shared Agent Representations

Frances Brazier^a Maarten van Steen^b Niek Wijngaards^a

^a Intelligent Interactive Distributed Systems, Vrije Universiteit
Amsterdam, de Boelelaan 1081a, 1081 HV Amsterdam

^b Computer Systems, Vrije Universiteit Amsterdam, de Boelelaan
1081a, 1081 HV Amsterdam

Reference. Brazier, F.M.T., van Steen, M., and Wijngaards, N.J.E. (2001). Distributed Shared Agent Representations. In: V. Marik, O. Stepankova, H. Krautwurmova, and J-P Briot (eds), *Proceedings of the Adaptability and Embodiment Using Multi-Agent Systems, AEMAS 2001 Workshop*, Prague, July 7th, 2001, pp. 230 - 237.

Extended Abstract

1

Different definitions of agents use concepts such as autonomy, pro-activity, reactivity, social abilities, and intentional models. Very few (if any), however, refer to an agent's external representation: the part of an agent that can be observed.

This representation is the visible part of an agent: (part of) its embodiment. This external representation may be limited to an agent's identity (e.g., name of owner or an IP address of the host on which it resides). It may also contain extensive public information about an agent (e.g., its profile), or it may even be a graphical figure (e.g., an avatar) representing an agent. The environment of an agent includes objects, and representations of itself and other agents. The distinction between an agent's external representation and an agent's internal processes and knowledge makes it possible to consider new ways to implement large scale agent systems.

In the multi-agent system community, large multi-agent systems are considered to consist of hundreds of agents, not thousands nor millions. As an example, consider the claim that Auctionbot is scalable, which is supported by an experiment with only 90 agents [3]. In the near future, however, we expect that multi-agent systems will need to be able to scale (in terms of the number of agents and available resources) to much larger populations. This almost immediately without noticeable loss of performance, or considerable increase in administrative complexity [1].

An agent can be seen as a (multi-threaded) process with internal knowledge, and an external representation. Replication of an entire agent is not an obvious option for the realisation of scalable systems (running processes in parallel on different machines will seldom be synchronous). In some cases cloning an agent

¹To appear in Proceedings of BNAIC2001.

may be a viable alternative to replication, but this is clearly application dependent and outside the scope of this paper.

Replication of public representations of agents, i.e. distributed shared agent representations, raises a number of issues regarding agents, and supportive middleware. These issues are related to policies (and/or strategies) for replication, accessibility, authority, and awareness. By combining choices for each of these issues, more and more complex situations may arise, for instance, whether an agents' public representation is replicated as soon as as soon as another agent wishes to observe it (i.e., many replicas), all agents may change a specific part of the public information of an agent, or an agent is aware of all changes to its public information. The middleware needed to support these combinations of policies, involves more expensive mechanisms (in terms of communication between entities in the multi-agent system) than a situation in which only the agent itself may modify its public representation. The question may even arise whether the more complex case scales well.

AgentScape is currently being designed to support the design and development of worldwide distributed, scalable, secure, and extensible agent systems. It aims to provide support in two ways. First, support is provided on the level of a basic agent operating system. Second, support is provided by services, such as location and directory services, automated creation of agents, and management of agents, objects, locations and groups. AgentScape provides basic building blocks needed to build such systems.

An agent in AgentScape consists of an external, visible part and an internal, invisible part. The external visible part of an agent may be observed by other agents, and contains the public representation of the agent. The internal (invisible) part of an agent includes local information, its process, data and/or knowledge. The visible part of an agent may be replicated, the invisible part of an agent is not replicated. Any amount of information may be included in the public representation of an agent; e.g., the visible part of an agent may be (nearly) empty. The public representation of an agent is implemented as a distributed shared GLOBE object [2].

The challenge is to design an environment in which replication of public representations of agents is possible and effective. Which replication strategies are most useful and applicable, which accessibility, authorisation and awareness policies are needed, are clearly, as yet, unanswered. Further research is clearly required!

References

- [1] B. Neuman. "Scale in Distributed Systems." In T. Casavant and M. Singhal, (eds.), *Readings in Distributed Computing Systems*, pp. 463–489. IEEE Computer Society Press, Los Alamitos, CA., 1994.
- [2] M. van Steen, P. Homburg, and A. Tanenbaum. "Globe: A Wide-Area Distributed System." *IEEE Concurrency*, 7(1):70–78, Jan. 1999.

- [3] P. R. Wurman, M. P. Wellman, and W. E. Walsh. “The Michigan Internet AuctionBot: A configurable auction server for human and software agents.” In K. P. Sycara and M. Wooldridge, (eds.), *Proc. Proceedings of the 2nd International Conference on Autonomous Agents (Agents’98)*, pp. 301–308, New York, 9–13, 1998. ACM Press.