

Nature-Inspired Multi-Agent Systems in the Task Assignment Domain

Steven de Jong, Karl Tuyls, Ronald Westra
Maastricht ICT Competence Center, Universiteit Maastricht,
P.O. Box 616, 6200 MD, Maastricht, The Netherlands

{steven.dejong, k.tuyls}@cs.unimaas.nl, westra@math.unimaas.nl

1. BACKGROUND

Task assignment problems are defined as problems in which multiple agents must efficiently share the load of executing and completing a certain set of tasks [REF]. *Resource distribution* is a task assignment problem in which multiple agents need to collect and deliver resources at pre-defined locations, while minimizing average waiting and delivery times for resources and clients. Research in the domain of resource distribution has a high industrial valorization; real-world examples include warehouses, airports, and hospitals. Because of this, there is increasing interest in robotic Multi-Agent Systems (MAS) that perform resource distribution tasks.

Agents situated in a resource distribution system have to (i) cooperate in an environment that is complex and highly dynamic, (ii) process unexpected requests (i.e., it is not known in advance where a resource should be picked up), and (iii) maintain efficient performance while taking into account possible failures of individual agents. As a consequence, any agent system designed for resource distribution tasks should be *scalable* with respect to the number of resources and collaborating robots, *adaptive* in order to handle unexpected requests and events, and *robust* to be able to deal with any kind of failure. Obviously, these properties have to be achieved while maintaining *functionality*.

Traditionally, resource distribution problems are countered by applying multi-agent planning (e.g., [4]). Recently, we can observe a shift towards *nature-inspired systems* (e.g., [5, 6, 7, 8, 9]). Here, many authors propose to use *ant systems* [10] and *swarm intelligence* [9, 11] in combination with pheromone-based and/or potential-field-based techniques, since these lead to reasonable scalability, flexibility and functionality, even on real robots (e.g., [8]).

In related work that suggests using nature-inspired approaches in resource distribution, we can clearly identify two major (closely related) shortcomings. First, systems that depend purely on emergent functionality are hard to control. We will give two examples here: (i) the relation between (individual) local behavior and the emergence of (collective) global behavior is often hard to predict, and (ii) nature-inspired systems tend to get stuck in local optima. Second, purely nature-inspired approaches often lack the *pragmatism* required for real-world applications, which simply should work. To deal with possible side effects of emergence, we should not exclude the possibility to introduce more traditional concepts into a nature-inspired approach.

NiSIS Mallorca Brainstorming Meeting 8–9 June, 2006 in Palma de Mallorca, Spain. This research has previously been addressed in [1, 2, 3]. Please consult these papers for a more complete overview of references.

2. OUR WORK

In order to overcome the problems observed in related work, we examine how design choices can be made that enable sufficient (or even optimal) conformance of the system to the four requirements: functionality, scalability, adaptivity and robustness. For example, to ensure that the system is scalable, its algorithmic complexity must be studied, and to ensure that it is functional, we must first determine what ‘being functional’ means in our context, and then investigate whether we can build a system that is indeed functional in this way.

Having performed this examination, we propose a technique which combines a robust mechanism found in nature with the pragmatism needed for real-world applications. It consists of (i) an *intelligent* environment incorporating both potential-field theory and straight-forward path planning, and (ii) a group of simple robots that respond only to local information. An example environment is shown in Figure 1. In summary:

- The environment is divided into a coarse *grid* which facilitates straightforward calculations and allows us to abstract out some of the physical difficulties of the task at hand (e.g., actually picking up resources).
- A simulated potential field is used to construct an *expected distribution* of future resources.
- Idle robots are attracted to the potential field and repelled from each other. They spread over the environment in a way that can be expected to decrease the average distances these robots have to travel when new resources appear.
- Busy robots are guided by traditional path (re-)planning, because this makes the system less sensitive to local optima and increases both robustness and performance, as shown in [3].

In our work, we introduce a quantitative measure for functionality in embodied MAS, i.e., the *fairness* of a system. Fairness enforces agents to be socially engaged towards each other and clients of the system. It can be studied using the fundamental ideas of Evolutionary Game Theory (EGT) [12]. It is known from this field and sociology that in many decision-making and strategic settings, people do not behave like self-interested rational agents as depicted in neoclassical economics and traditional game theory; instead, individuals usually have an inequality aversion. Typically, this leads to the emergence of *altruism* in ultimatum and public good games [12, 13], of which resource distribution can be considered an instance. Support for these ideas comes from anthropological literature, describing how *Homo sapiens* evolved in small hunter-gatherer groups [12]. Such societies had no centralized structure




Figure 1: An example grid environment with various walls. Potentials are visualized by color and denoted in each cell, below the cell's coordinate. Gray cells lie outside the environment. Robots are visualized using circles (outlined for idle robots, filled for robots that are going to a resource). Resources are visualized using small black boxes. It can be seen that some resources are currently held by a robot (e.g., at (4,4)).

of governance, so the enforcement of norms depended on the voluntary participation of agents. We argue that functionality of a system is typically addressed based on the rationality assumptions of classical game theory, showing optimality criteria that do not match with human expectations and decision making. We illustrated this by an experiment with human actors [3], which shows that in a real society, agents do not only care about their own utility but also about how this utility compares to the utility of others.

3. RESEARCH AGENDA

For future research, we will extend our current research along two nature-inspired lines, viz. one related to biology-inspired techniques and one related to evolutionary game theory. Both lines will be briefly illustrated here.

Biology-inspired techniques. Our research is divided in two distinct phases. In the first phase, we allow a central control approach. In the second phase, we will study how we can further increase adaptivity and robustness of the system by introducing distributed control, for example by introducing sensor networks. We believe that it is important to clearly divide both complex tasks, since the large number of different effects playing a role in both of these tasks could make results difficult to interpret.

Evolutionary game theory. The notion of fairness we found in our experiments, cannot only be used as a measure for existing systems, but can also be applied explicitly in the design of a new system. Currently, we are investigating how to develop an explicitly fair resource distribution system. In [3], we found that humans tend to select the *center of gravity* in simplified scenarios with only one agent and a limited number of locations where resources are offered. Generalizing this result to a situation with more agents can be done, for example, by partitioning the environment into regions (as many regions as there are agents) and determining the center of gravity for each of these regions (and thus, determining the fair position for an agent within the region). The size and shape

of the regions can emerge during the simulation, using the actual assignments of agents to resources our planner performs. Obviously, such a system is especially needed when agents are regularly idle – if they are constantly busy, they will not have time to position themselves at the required center of gravity.

Acknowledgments. This research reported here is being carried out under the European research grant FP6-013569 NiSIS (Nature-inspired Smart Information Systems), running until March, 2008. It is part of the Interactive Collaborative Information Systems (ICIS) project, supported by the Dutch Ministry of Economic Affairs, grant nr: BSIK03024, and is partially funded by the Breedtestrategie programme of the Universiteit Maastricht.

4. REFERENCES

- [1] Steven de Jong, Karl Tuyls, and Ida Sprinkhuizen-Kuyper. Nature-Inspired Multi-Agent Coordination in Task Assignment Problems. In Ann Nowe, Maarten Peeters, and Katja Verbeeck, editors, *Proceedings of the 6th European Symposium on Adaptive Learning Agents and MAS (ALAMAS)*, 2006.
- [2] Steven de Jong, Karl Tuyls, Tsuyoshi Hashimoto, and Hiroyuki Iida. Scalable Potential-Field Multi-Agent Coordination in Resource Distribution Tasks. In Liviu Panait, Sandip Sen, and Eduardo Alonso, editors, *Proceedings of the AAMAS 2006 Workshop on Adaptation and Learning in Autonomous Agents and Multiagent Systems*, Hakodate, Hokkaido, Japan, May 8th 2006.
- [3] Steven de Jong, Karl Tuyls, Benjamin Torben-Nielsen, and Ida Sprinkhuizen-Kuyper. Towards Fair Resource Distribution by Socially Engaged Agents. In *Proceedings of PPSN 2006 (submitted)*, 2006.
- [4] Dani Goldberg, Vincent Cicirello, M Bernardine Dias, Reid Simmons, Stephen Smith, and Anthony (Tony) Stentz. Market-Based Multi-Robot Planning in a Distributed Layered Architecture. In *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings from the 2003 International Workshop on Multi-Robot Systems*, volume 2, pages 27–38. Kluwer Academic Publishers, 2003.
- [5] Danny Weyns, Nelis Boucke, Tom Holvoet, and Wannes Schols. Gradient Field-Based Task Assignment in an AGV Transportation System. In *EUMAS 2005 - Proceedings of the Third European Workshop on Multi-Agent Systems, Brussels, Belgium, December 7-8, 2005*, pages 447–458, 2005.
- [6] Liviu Panait and Sean Luke. A Pheromone-Based Utility Model for Collaborative Foraging. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 36–43, Washington, DC, USA, 2004. IEEE Computer Society. ISBN 1-58113-864-4.
- [7] H. Van Dyke Parunak, S. Brueckner, and J. Sauter. Synthetic Pheromone Mechanisms for Coordination of Unmanned Vehicles. In *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems*, pages 448–450, Bologna, Italy, 2002. ACM Press, New York, NY.
- [8] David Payton, Mike Daily, Regina Estowski, Mike Howard, and Craig Lee. Pheromone Robotics. *Auton. Robots*, 11(3):319–324, 2001.
- [9] O.E. Holland and C. Melhuish. Stigmergy, Self-organisation, and Sorting in Collective Robotics. *Artificial Life*, 5(2):173–202, 1999.
- [10] Shu-Chuan Chu, John F. Roddick, and Jeng-Shyang Pan. Ant colony with communication strategies. *Information Sciences*, 167:63–76, 2004.
- [11] S. A. Brueckner, G. Di Marzo Serugendo, A. Karageorgos, and R. Nagpal (eds). *Engineering Self-Organising Systems: Methodologies and Applications*. In *LNCS Postproceedings of ESOA 2004*. Springer LNAI 3464, 2005.
- [12] H. Gintis. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Number ISBN: 0-691-00943-0. Princeton University Press, 2001.
- [13] Jorgen W. Weibull. *Evolutionary Game Theory*. MIT Press, 1996.