

# Evolving Real-Time Local Agent Control for Large-Scale MAS

Thomas Wagner  
Computer Science Department  
University of Maine  
wagner@umcs.maine.edu

Victor Lesser  
Computer Science Department  
University of Massachusetts  
lesser@cs.umass.edu

## ABSTRACT

Control for agents situated in multi-agent systems is a complex problem. This is particularly true in hard, open, dynamic environments where resource, privacy, bandwidth, and computational limitations impose restrictions on the type of information that agents may share and the control problem solving options available to agents. The *MQ* or *motivational quantities* framework addresses these issues by evaluating candidate tasks based on the agent's organizational context and by framing control as a local agent optimization problem that approximates the global problem through the use of state and preference.

## 1. MQ OVERVIEW

Our objective is to create *open, large-scale*, information and computational systems that are flexible, adaptable, robust, persistent, and autonomous. Now, consider the implications of this. Openness means that agents may interact freely, come and go from the network, and that the entire problem solving environment is dynamic. Openness thus often acts to thwart agent technologies that rely on detailed predictability or static properties of the problem space. In our work, we take the view that openness leads to a requirement for real-time agent control problem solving so that agents can respond to change and unexpected outcomes online.

Moving the scale of multi-agent systems from small groups to large groups, e.g., tens of thousands, throws two other problems into the mix: *increased interaction overhead* and *social complexity*. The term interaction overhead denotes the increase in communication between agents required to detect interactions in their problem solving and to coordinate their activities, i.e., it denotes the sheer volume of message traffic and problem solving required to evaluate the messages. This is being dealt with by imposing organizational structure on the agents so that they do not all communicate and by creating coordination and communication technologies that are adjustable [1, 4, 15, 5]. The other issue is social complexity and we do not mean social complexity in the human sense. Or rather, the goal of this research is not to study social complexity in human organizations [13] per se as our work in agent control has very specific task-centered properties. When agents are situated in a large open environment, and organizational structure is imposed upon them,

they have different organizational objectives and they must reason about how their problem solving relates to satisfying their multiple, and possibly conflicting, organizational objectives.

This research focuses on exactly this problem – how agents in large-scale open environments reason about their organizational context and make appropriate choices about which actions to perform and how to go about performing them. It is important to emphasize that this research pertains to complex problem solving agents, e.g., the BIG Information Gathering Agent [11], where the agents are situated in an environment, able to sense and effect, and have explicit representations of candidate tasks and explicit representations of different ways to go about performing the tasks. Additionally, tasks are quantified or have different performance characteristics and, following in the thread of complex problem solving there are relationships between the tasks. The implications are that tasks cannot be reasoned about independently and that the value or utility of particular tasks differs depending on the context. We call the process of reasoning about which tasks to perform, when, with what resources, how or in what fashion, and with whom to coordinate, the *local agent control* problem. The term “local” is used in this expression because agency, as we use it, denotes an autonomous distributed problem solving entity. In our work, there is no global picture of all activities being carried out by all agents nor are the agents situated in specialized, tightly coupled environments like Tambe's teams [15] or robotic soccer [17].

We view local agent control in this context as a real-time action-selection-sequencing problem where an agent has  $n$  candidate tasks and alternative different ways to perform the tasks. Tasks have deadlines and other constraints as well as different performance properties, e.g., consuming different resources or producing results of varying quality. Control in this context is an optimization problem where different solutions are possible and they have different degrees of utility.

Historically in our work this class of control problem has been dealt with using the TÆMS task modeling framework [2, 10], GPGP coordination [2], and Design-to-Criteria (DTC) real-time agent scheduling [14, 19, 18, 21, 7]. Using these tools, an individual agent for use in a multi-agent environment is constructed by coupling a domain expert or planner with GPGP and DTC. In this model, the domain expert's function is to perform domain problem solving and to translate its internal representations into TÆMS for control problem solving by the coordination (GPGP) and trade-off/scheduling (DTC) experts. GPGP and DTC then work together to guide the actions of the individual agent and to coordinate the activities of the agent with the other agents in the network. This is the approach used in the BIG information gathering agent [12, 11], the Intelligent Home project (IHome) [9], the DARPA ANTS real-time agent sensor network for vehicle tracking [7], and others [22]. Though

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AGENTS'01, May 28-June 1, 2001, Montréal, Quebec, Canada.  
Copyright 2001 ACM 1-58113-326-X/01/0005 ...\$5.00.

in some of these applications GPGP is replaced by other communication machinery that forms commitments between agents about which tasks will be performed and when. In all of these applications, DTC or its predecessor, Design-to-Time [8], is the oracle that guides and constrains the communication and commitment formation processes.

TÆMS and DTC are mature research artifacts and have been successfully reused in many applications (DTC since 1995). However, TÆMS is not suited to addressing the situational complexity that arises when agents are deployed in larger groups or in open environments. One of the fundamental limitations of TÆMS is that it is a static representation of an agent's problem solving process at a given instant in time. It is, in essence, a snapshot of the options available to the agent and a snapshot of their characteristics. In our applications, generally, when the situation changes and the characteristics of tasks (used to determine utility) change, the problem solver must adjust the performance profiles and emit a new TÆMS task structure. Another limitation is that in TÆMS, action performance produces *quality* which then propagates throughout the entire graph-like structure in ways that is intended to model distributed problem solving as in a distributed interpretation problem [6]. The formal details of TÆMS are in [3]. While this view is appropriate for reasoning about interrelated domain problem solving activities at a detailed level, it is not readily used to model concepts like tasks that contribute to one organizational objective while being detrimental to another. TÆMS also does not adequately support concepts like the value of forming a commitment with another agent or the penalty for decommitting from an activity once a commitment is formed.

To address these limitations, we have developed a new framework for representing tasks and actions at a different level of abstraction. The framework, called the *motivational quantities* ( $MQ$ ) [16, 20] framework, uses state to achieve "automatic" changes in task valuation or utility (unlike the static view taken in TÆMS). The  $MQ$  framework also describes tasks in many different attribute dimensions so that we can model tasks contributing to, or detracting from, different objectives to different degrees. While control at the TÆMS level pertains to detailed evaluation of domain problem solving activities of an agent, control at the  $MQ$  level pertains to high-level valuation of candidate tasks based on an understanding of the relationship between tasks and organizational objectives. In other words, in the  $MQ$  framework, task value is determined not only by the intrinsic properties of tasks, but by the benefits and costs of the intrinsic properties as determined by the agent's current organizational situation. From another view, there is an intermediate evaluation step in the control process whereas such processes typically focus on intrinsic value rather than contextually interpreted value. While we have ideas about how to combine and interface [16] the two levels, integration is clearly unnecessary for many applications.

A preliminary version of the  $MQ$  framework was presented in [20]. A specification of the current model is located at [23].

## 2. REFERENCES

- [1] K. Decker, A. Pannu, K. Sycara, and M. Williamson. Designing behaviors for information agents. In *Proc. of the 1st Intl. Conf. on Autonomous Agents*, February 1997.
- [2] Keith Decker and Jinjiang Li. Coordinated hospital patient scheduling. In *Proc. of the 3rd Intl. Conf. on Multi-Agent Systems (ICMAS98)*, 1998.
- [3] Keith S. Decker. *Environment Centered Analysis and Design of Coordination Mechanisms*. PhD thesis, Univ. of Massachusetts, 1995.
- [4] Keith S. Decker and Victor R. Lesser. Designing a family of coordination algorithms. In *Proc. of the 1st Intl. Conf. on Multi-Agent Systems*, AAAI Press, 1995.
- [5] C. Dellarocas and M. Klein. An experimental evaluation of domain-independent fault handling services in open multi-agent systems. In *Proc. of the 5th Intl. Conf. on Multi-Agent Systems*, 2000.
- [6] Edmund H. Durfee and Victor R. Lesser. Using partial global plans to coordinate distributed problem solvers. In *Proc. of the Tenth Intl. Joint Conf. on AI*, August 1987.
- [7] Regis Vincent et al. Distributed Sensor Network for Real Time Tracking. In *Proc. of Autonomous Agents*, 2001. To appear.
- [8] Alan J. Garvey. *Design-to-Time Real-Time Scheduling*. PhD thesis, Univ. of Massachusetts, February 1996.
- [9] Victor Lesser, Michael Atighetchi, Bryan Horling, Brett Benyo, Anita Raja, Regis Vincent, Thomas Wagner, Ping Xuan, and Shelley XQ. Zhang. A Multi-Agent System for Intelligent Environment Control. In *Proc. of the 3rd Intl. Conf. on Autonomous Agents*, 1999.
- [10] Victor Lesser, Bryan Horling, and et al. The TÆMS whitepaper / evolving specification. <http://mas.cs.umass.edu/research/taems/white>.
- [11] Victor Lesser, Bryan Horling, Frank Klassner, Anita Raja, Thomas Wagner, and Shelley XQ. Zhang. BIG: An agent for resource-bounded information gathering and decision making. *Artificial Intelligence*, 118(1-2):197-244, May 2000. Elsevier Science Publishing.
- [12] Victor Lesser, Bryan Horling, Anita Raja, Thomas Wagner, and Shelley XQ. Zhang. Sophisticated Information Gathering in a Marketplace of Information Providers. *IEEE Internet Computing*, 4(2):49-58, Mar/Apr 2000.
- [13] Michael J. Prietula, Kathleen M. Carley, and Les Gasser. A Computational Approach to Organizations and Organizing. In Michael J. Prietula, Kathleen M. Carley, and Les Gasser, editors, *Simulating Organizations: Computational Models of Institutions and Groups*, pages xiv-xix. AAAI Press / MIT Press, 1998.
- [14] Anita Raja, Victor Lesser, and Thomas Wagner. Toward Robust Agent Control in Open Environments. In *Proc. of the 4th Intl. Conf. on Autonomous Agents*, 2000.
- [15] Milind Tambe. Agent Architectures for Flexible, Practical Teamwork. In *Proc. of the 14th National Conf. on AI*, July 1997.
- [16] Thomas A. Wagner. *Toward Quantified Control for Organizationally Situated Agents*. PhD thesis, Univ. of Massachusetts, February 2000.
- [17] Manuela Veloso, Peter Stone, and Kwun Han. The CMUnited-97 robotic soccer team: Perception and multiagent control. In *Proc. of the 2nd Intl. Conf. on Autonomous Agents (Agents98)*, pages 78-85, 1998.
- [18] Thomas Wagner, Brett Benyo, Victor Lesser, and Ping Xuan. Investigating Interactions Between Agent Conversations and Agent Control Components. In Frank Dignum and Mark Greaves, editors, *Issues in Agent Communication*, Lecture Notes in AI, Springer-Verlag, Berlin, 2000.
- [19] Thomas Wagner, Alan Garvey, and Victor Lesser. Criteria-Directed Heuristic Task Scheduling. *Intl. Journal of Approximate Reasoning, Special Issue on Scheduling*, 19(1-2):91-118, 1998. A version also available as UMASS CS TR-97-59.
- [20] Thomas Wagner and Victor Lesser. Relating quantified motivations for organizationally situated agents. In N.R. Jennings and Y. Lespérance, editors, *Intelligent Agents VI (Proc. of ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000.
- [21] Thomas Wagner and Victor Lesser. Design-to-Criteria Scheduling: Real-Time Agent Control. In Thomas Wagner and Omer Rana, editors, *To appear in Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems*, LNCS. Springer-Verlag, 2001. Also appears in the 2000 AAAI Spring Symposium on Real-Time Systems and a version is available as Univ. of Massachusetts Computer Science Technical Report TR-99-58.
- [22] Thomas Wagner, John Phelps, Yuhui Qian, Erik Albert, and Glen Beane. A modified architecture for constructing real-time information gathering agents. In *Proc. of Agent Oriented Information Systems*, 2001. To appear.
- [23] Thomas Wagner <http://www.umcs.maine.edu/~wagner/mq>