In *Proc. of the AAAI Spring Symposium on Bridging the Multi-Agent and Multi-Robotic Research Gap* pages 1-3, Palo Alto, California, March, 2004

Are (explicit) multi-robot coordination and multi-agent coordination really so different?

Brian P. Gerkey
Robotics Laboratory
Stanford University
Stanford, CA 94305-9019, USA
gerkey@robotics.stanford.edu

Maja J Matarić
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781, USA
mataric@cs.usc.edu

Largely because of significant qualitative differences between the respective systems under study, the multi-agent and multi-robot research communities have each developed their own methods for perception, reasoning, and action in individual agents/robots. In particular, the multi-robot community has historically studied both implicit and explicit coordination techniques. *Implicit coordination* techniques employ dynamics of interaction among the robots and the environment in order to achieve the desired collective performance, often in the form of designed emergent behavior. These methods, while often very elegant and efficient, have so far defied general analysis, but show great promise in particular for large-scale teams of simple individuals, and are being studied actively in robotics as well as in several other fields (including mathematics, artificial life, etc.). *Explicit coordination* techniques, in comparison, deal with comparatively more sophisticated agents/robots, and employ intentional communication and collaboration methods much like those employed in multi-agent systems. Therefore, at the level of explicit coordination among multiple individuals, the differences between techniques used in multi-agent systems (MAS) and those used in multi-robot systems (MRS) are in fact very few. This is not to say that MAS and MRS are equivalent in any fundamental way, but rather that although robotics researchers employ sophisticated specialized techniques of various sorts (e.g., control-theoretic (Hao, Laxton, Agrawal, Lee & Benson 2003, Pereira, Das & Kumar 2003), probabilistic (Riley & Veloso 2002, Thomas Röfer and Matthias Jüngel 2003)) when designing single-robot control systems, they have so far tended to use techniques that are already well-known in the agent community when designing explicitly coordinated MRS. The discussion in the rest of this paper deals with explicitly coordinated MRS, so MRS can heretofore be assumed to refer to such systems and to exclude implicit MRS.

In our previous work (Gerkey & Matarić 2003*a*), we examined the similarity among existing MRS in some depth, focusing on the problem of multi-robot task allocation (MRTA). We showed that many of the MRTA architectures that can be found in the literature are in fact solving well-understood optimization problems (e.g., the Online Assignment Problem (Kalyanasundaram & Pruhs 1993)) using well-understood techniques (e.g., the canonical Greedy Algorithm (Ahuja, Magnanti & Orlin 1993)). We also offered a theoretically sound, non-economic explanation for the success of market-based task allocation mechanisms (e.g., (Dias & Stentz 2001), (Gerkey & Matarić 2002)) based on well-known results from linear programming (Gale 1960).

In a recent extension to our analysis (Gerkey & Matarić 2003*b*), we developed a taxonomy of MRTA problems, dividing the space along three axes:

- **single-task robots (ST)** vs. **multi-task robots (MT)**: ST means that each robot is capable of executing as most one task at a time, while MT means that some robots can execute multiple tasks simultaneously.
- **single-robot tasks (SR)** vs. **multi-robot tasks (MR)**: SR means that each task requires exactly one robot to achieve it, while MR means that some tasks can require multiple robots.
- **instantaneous assignment (IA)** vs. **time-extended assignment (TA)**: IA means that the available information concerning the robots, the tasks, and the environment permits only an instantaneous allocation of tasks to robots, with no planning for future allocations. TA means that more information is available, such as the set of all tasks that will need to be assigned, or a model of how tasks are expected to arrive over time.

We had two goals in mind for this taxonomy; (i) to show how various MRTA problems can be positioned in the resulting problem space; and (ii) to explain how organizational theory (e.g., operations research, scheduling) relates to those problems and to proposed solutions from the robotics literature. In some cases, we were able to construct provably optimal algorithms, as well as give bounds for the solution quality that can be expected from the suboptimal algorithms currently in use in MRS research. In other cases, only approximate solutions are available, and for some difficult MRTA problems, there do not currently exist good approximations.

Interestingly, neither in our analysis, nor in the existing MRTA architectures that we studied, were significant robot-specific decisions or assumptions made regarding coordination. That is to say, modulo implementation details such as communication timeouts, these *MRS* coordination mechanisms could fairly be described as *MAS* coordination mechanisms. Nonetheless, most of the MRS architectures we analyzed have been validated with physical robots engaged in various coordinated tasks. As a result, we know that at least *some* techniques, especially simple ones, that are used for coordinating MAS are also likely to achieve some success when applied to MRS. These techniques were not necessarily developed *by* the MAS community, and in fact are often borrowed or adapted from other fields, such as operations research and economics. Owing to the continual increase in available computing resources for physical robots, this class of MAS-style algorithms that can be effectively used in MRS has grown significantly to include, for example, the use of Markov decision processes (MDPs) (Kaelbling, Littman & Cassandra 1998). Recent work in factoring MDPs into tractable chunks (Guestrin, Koller & Parr 2001) has allowed for their use on line for coordinating the actions of (small) teams of robots (Rosencrantz, Gordon & Thrun 2003).

However, it remains an open question as to how much benefit can be derived from using sophisticated coordination methods in MRS, because of an important underlying issue, which we suggest is one of the primary challenges facing MRS (and MAS): *utility*. How is the utility of a given course of action for a given robot or group of robots to be decided? This question is difficult to answer for a single robot and even harder for a MRS. However, it must be answered, because most if not all coordination approaches rely on some form of utility, whether referred to by that name (Chaimowicz, Campos & Kumar 2002, Gerkey & Matarić 2003*a*), or as eligibility (Werger & Matarić 2001), capability (Smith 1980), fitness (Gerkey & Matarić 2002), cost (Botelho & Alami 1999, Dias & Stentz 2002), or reward (Bererton, Gordon & Thrun 2003).

In any case, since coordination is achieved by maximizing utility (or, equivalently, minimizing cost), the utility measure *must* account for all state information that is relevant to the task. All information that affects task performance but is not captured in the utility measure is captured in what economists refer to as *externalities*, the effects of which can be disastrous (Simon 2001). In fact, it is usually the case that the quantity being optimized, utility, is not a direct measure of task performance, nor is it necessarily even strongly correlated with task performance. As a result, an "optimal" coordination solution, which maximizes utility, may not actually produce optimal (or even good) system performance. A common externality encountered in MRS is *physical interference*, which is often ignored or only crudely modeled when estimating utilities but can have complex and unpredictable effects and may easily dominate task performance (Goldberg & Matarić 1997). Simple coordination approaches like those used in most MRTA research to date cannot account for interference effects, because they assume (often implicitly) that the robots' utilities for tasks are independent and that the total system metric is a linear combination of these values. Alternatively, MDPs can account for interference effects, but at the expense of complicating the model and adding constraints among otherwise weakly-connected components of the model. This last point is important because current methods for efficiently solving MDPs rely on decomposing the model into small, fairly independent, chunks, and then integrating the results, often using a linear program (which again implies a linear combination of utilities).

We suggest that a vital area for future research is the principled derivation of utility values, both for MAS and MRS. Without a method for defining meaningful utility values for situated systems, better and more efficient methods for coordination through utility optimization will have a limited practical impact on system performance.

References

Ahuja, R. K., Magnanti, T. L. & Orlin, J. B. (1993), *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, Upper Saddle River, New Jersey.

Bererton, C., Gordon, G. & Thrun, S. (2003), Auction Mechanism Design for Multi-Robot Coordination, *in* 'Proc. of Advances in Neural Information Processing Systems (NIPS)', Vancouver, Canada. To appear.

Botelho, S. C. & Alami, R. (1999), M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement, *in* 'Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)', Detroit, Michigan, pp. 1234–1239.

Chaimowicz, L., Campos, M. F. M. & Kumar, V. (2002), Dynamic Role Assignment for Cooperative Robots, *in* 'Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)', Washington, DC, pp. 293–298.

Dias, M. B. & Stentz, A. (2001), A Market Approach to Multirobot Coordination, Technical Report CMU-RI-TR-01-26, The Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania.

Dias, M. B. & Stentz, A. (2002), Opportunistic Optimization for Market-Based Multirobot Control, *in* 'Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)', Lausanne, Switzerland, pp. 2714–2720.

Gale, D. (1960), *The Theory of Linear Economic Models*, McGraw-Hill Book Company, Inc., New York.

Gerkey, B. P. & Matarić, M. J. (2002), 'Sold!: Auction methods for multi-robot coordination', *IEEE Transactions on Robotics and Automation* **18**(5), 758–768.

Gerkey, B. P. & Matarić, M. J. (2003*a*), A Framework for Studying Multi-Robot Task Allocation, *in* A. Schultz et al., eds, 'Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II', Kluwer Academic Publishers, the Netherlands, pp. 15–26.

Gerkey, B. P. & Matarić, M. J. (2003*b*), A formal framework for the study of task allocation in multi-robot systems, Technical Report CRES-03-013, Center for Robotics and Embedded Systems, School of Engineering, University of Southern California.

Goldberg, D. & Matarić, M. J. (1997), Interference as a tool for designing and evaluating multi-robot controllers, *in* 'Proc. of the Natl. Conf. on Artificial Intelligence (AAAI)', Providence, Rhode Island, pp. 637–642.

Guestrin, C., Koller, D. & Parr, R. (2001), Multiagent Planning with Factored MDPs, *in* 'Proc. of Advances in Neural Information Processing Systems (NIPS)', Vancouver, Canada, pp. 1523–1530.

Hao, Y., Laxton, B., Agrawal, S. K., Lee, E. & Benson, E. (2003), Planning and Control of UGV Formations in a Dynamic Environment: A Practical Framework with Experiments, *in* 'Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)', Taipei, Taiwan, pp. 1209–1214.

Kaelbling, L. P., Littman, M. L. & Cassandra, A. R. (1998), 'Planning and Acting in Partially Observable Stochastic Domains', *Artificial Intelligence* **101**(1–2), 99–134.

Kalyanasundaram, B. & Pruhs, K. (1993), 'Online Weighted Matching', *J. of Algorithms* **14**, 478–488.

Pereira, G. A. S., Das, A. K. & Kumar, V. (2003), Decentralized Motion Planning for Multiple Robots Subject to Sensing and Communication Constraints, *in* A. Schultz et al., eds, 'Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II', Kluwer Academic Publishers, the Netherlands, pp. 267–278.

Riley, P. & Veloso, M. (2002), Planning for distributed execution through use of probabilistic opponent models, *in* 'Proc. of the Intl. Conf. on Artificial Intelligence Planning Systems', Toulouse, France, pp. 72–81.

Rosencrantz, M., Gordon, G. & Thrun, S. (2003), Locating Moving Entities in Indoor Environments with Teams of Mobile Robots, *in* 'Proc. of Intl. Conf. on Autonomous Agents and Multi Agent Systems', Melbourne, Australia, pp. 233–240.

Simon, H. A. (2001), *The Sciences of the Artificial*, 3$^{\mathrm{rd}}$ edn, MIT Press, Cambridge, Massachusetts.

Smith, R. G. (1980), 'The Contract Net Protocol', *IEEE Tranactions on Computers* **29**(12), 1104–1113.

Thomas Röfer and Matthias Jüngel (2003), Vision-Based Fast and Reactive Monte-Carlo Localization, *in* 'Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)', Taipei, Taiwan, pp. 856–861.

Werger, B. B. & Matarić, M. J. (2001), Broadcast of Local Eligibility for Multi-Target Observation, *in* L. E. Parker, G. Bekey & J. Barhen, eds, 'Distributed Autonomous Robotic Systems 4', Springer-Verlag, New York, pp. 347–356.