

Challenges for Multi-Agent Coordination Theory Based on Empirical Observations

Victor Lesser and Daniel Corkill
College of Information and Computer Sciences
University of Massachusetts Amherst

(An extended and revised version of a paper presented at AAMAS 2014 under the same title)

Abstract

Significant research progress and understanding about the nature of coordination has been made over the years. Development of the DCOP and DEC-MDP frameworks in the past decade has been especially important. Although these advances are very important for multi-agent coordination theory, they overlook a set of coordination behaviors and phenomena that have been observed empirically by many researchers since the early years of the field. The goal of this paper is to challenge researchers in multi-agent coordination to develop a comprehensive formal framework that explains these empirical observations.

Introduction

The study of coordination and cooperation among agents has been at the heart of the multi-agent field since its inception [Lesser81:Functionally, Davis83:Negotiation]. Coordination problems arise when: 1) an agent has a choice in its actions within some task, and the choice affects other agents' performance; 2) the order in which actions are carried out affects other agents' performance; or 3) the time at which actions are carried out affects other agents' performance. A coordination strategy involves choosing what actions to take, how to take them, when to take them, and by whom. It may also involve calculating and exchanging meta-information (e.g., state information about an agent's problem-solving, such as the next actions that are to be executed) among agents to facilitate coordination decisions. When there is no exchange of meta-information among agents, we generally call this an implicit or off-line form of coordination to differentiate it from an explicit and on-line form of coordination where meta-information is communicated.

Since this early work, significant research progress and understanding about the nature of coordination has been made [Durfee91:Partial, Lesser91:Retrospective Jennings93:Commitments, Tambe97:Towards, Yokoo98:Distributed, Lesser98:Reflections, Lesser04:Evolution]. Especially important has been the development of distributed constraint optimization (DCOP) [Yokoo91:Distributed] and decentralized Markov decision processes (DEC-MDPs) [Bernstein00:Complexity] frameworks over the last decade. These formal frameworks allow researchers to understand not only the inherent computational complexity of coordination problems, but also how to construct optimal or near-optimal coordination strategies for a wide variety of multi-agent applications. The DCOP framework is generally used for dynamic and on-line distributed coordination problems involving a single-shot control decision strategy for each

agent in the network. The word “dynamic” is used here to connote the idea that a new DCOP problem is constructed for each coordination cycle.¹ It applies to coordination problems where a greedy and incremental approach to coordination is an effective strategy. In contrast, the DEC-MDP framework has been used to create static (implicit), off-line coordination strategies involving a set of related sequential decisions for each agent where there is uncertainty over the outcome of each decision. It is generally applied to problems where there exists a model of the environment and agent task processing capabilities, and where there is a finite decision horizon.² The DEC-MDP framework has also been used to create an evolving, on-line coordination strategy for problems where a priori knowledge of the environment does not exist, and where there is an infinite horizon decision-making process. This latter approach is associated with work on multi-agent reinforcement learning.

The recent development of these two frameworks has emphasized how to take the specific structural properties of the different coordination problems such as [Becker04:Solving, Goldman04:Decentralized, Nair05:NETPOMDPS, Seuken08:FormalModels, Tarlow2010:HOP, Mostafa11:Compact, Witwicki11:Towards, Yeoh2013:Automated, Kumar11:Scalable; Pujol2013:Binary] to reduce computational and communication effort. Exciting challenges for each of these frameworks remain. For the DEC-MDP framework, there is the problem of computing off-line policies in realistic time for larger and more complex problems and the problem of making this framework useable in situations where the coordination problem is not completely static (describable off-line) but can vary to some degree on each coordination cycle [Yeoh2013:Automated]. For DCOPs, there are similar scaling issues in making them practical for larger and more complex applications in terms of reducing both run time computational and communication requirements. Many of these scaling approaches involve the use of approximate solutions to the coordination problems [Farinelli08:Decentralized]. There also has been some work on efficient solutions to a DCOP in situations where the DCOP problems from one coordination cycle to another are similar but not exactly the same [Macarthur10:Superstabilizing, Zivan10:Distributed].

These directions are very important for multi-agent coordination theory but overlook a set of coordination behaviors and phenomena that have been observed empirically by researchers since the early years of the field: 1) that structural interrelationships among agent activities inherent in the problem description are not necessarily indicative of the communication complexity necessary for effectively coordinating agents, 2) that implicit control without communication works so well, 3) that if you are willing to accept non-optimal solutions then coordination requirements can be dramatically reduced with often only a slight loss in performance, 4) that modifying local problem solving to make it more predictable or less responsive/opportunistic or decreasing the frequency of coordination sometimes improves agent coordination, 5) that a greedy and incremental approach to coordination can often lead to near-optimal solutions, 6) that sophisticated coordination strategies (in comparison to simpler approaches) are most effective only for limited classes of problems/task environments, and 6) that dynamic adaptation of a coordination strategy

¹ In principle, a DCOP can be used to represent a sequential set of decisions where there is no uncertainty associated with the outcome of each decision but often the combinatorics make this infeasible.

² The possibility of agents’ communicating meta-information among themselves increases significantly the computational difficulty of finding an optimal DEC-MDP coordination policy.

to the current state of the network problem solving can lead to more effective coordination. These behaviors have often been exploited by researchers building efficient heuristic coordination mechanisms, but rarely are they understood deeply or explained formally. There are, however, some exceptions [Decker93:Approach, Sen1998:Meeting], but they are limited to specific and narrowly defined coordination problems. Exploiting these phenomena usually requires taking a more statistical view of coordination behavior and taking into consideration the underlying distributed search process being coordinated. This is in contrast with current formal approaches that look for some explicit structural interaction pattern associated with a problem description that reduces computational complexity.

The goal of this paper is to challenge researchers in multi-agent coordination to develop a comprehensive formal framework that explains these empirical observations. A deeper, formal understanding of these phenomena could help researchers develop new and more efficient coordination strategies—possibly similar to how the study of phase transitions in NP-hard problems [Monasson99:Determining] opened up new perspectives to researchers studying computational complexity and search mechanisms.

Coordination and Communication

Most of the formal research in coordination theory uses the explicit structural patterns of interaction among agents to reduce the computational effort required to find an optimal strategy. These structural relationships are typically obtained from characteristics of the problem description. An early example of this approach was work on transition independent DEC-MDPs [Becker04:Solving], where actions taken in one agent do not affect the outcome of actions taken by another agent. However, we hypothesize that the existence of these structural relationships does not always indicate the communication requirements necessary for implementing an effective coordination strategy. An example of this was the observation by Mostafa [Mostafa11:Private] that, for at least one class of problems with structural interaction patterns that (on the surface) indicated that explicit communication of agent states would be advantageous, it was very hard to find specific problem instances of this class where the optimal coordination strategy actually required communication. A slightly different but related observation was made in the early work on solving DEC-MDPs by Xuan et al. [Xuan02:Multi-agent]. The approach they used first solved a centralized version of the coordination problem framed as a Multi-Agent Markov Decision Process (MMDP). In this MMDP solution, agents were aware of the state of other agents, which implied that the distributed implementation of each agent's policy required the agent to communicate its current state to other agents at each time step. Through analysis of this optimal centralized policy, it was shown that many of these communication actions were unnecessary, and that an optimal coordination policy could still be maintained in at least two-agent examples.

From our perspective, what was even more interesting occurred when approximations were introduced to this optimal policy derived from the MMDP solution by assuming, from a statistical perspective, what were the likely problem-solving states of the other agent given their local control policies. With these heuristics, they demonstrated that they could eliminate large amounts of communication with only a slight reduction in optimality (see Figure 1). More generally, the permissible orderings of local ordering of agent activities can often be exploited implicitly by the

coordination strategy to reduce the need for explicit coordination among agents. It is our hypothesis that, to the degree that there is more flexibility in how to organize local problem solving, it becomes more likely that the coordination strategy can find a combined ordering of local agent activities that reduces the need for explicit coordination among agents. From this perspective, the introduction of non-optimal local behavior, if done astutely, can present new options for finding combined agent activity orderings; thus, potentially reducing coordination overhead.



With a Slight Decrease in Optimality of Solution you can Dramatically Lessen the Amount of Communication Required

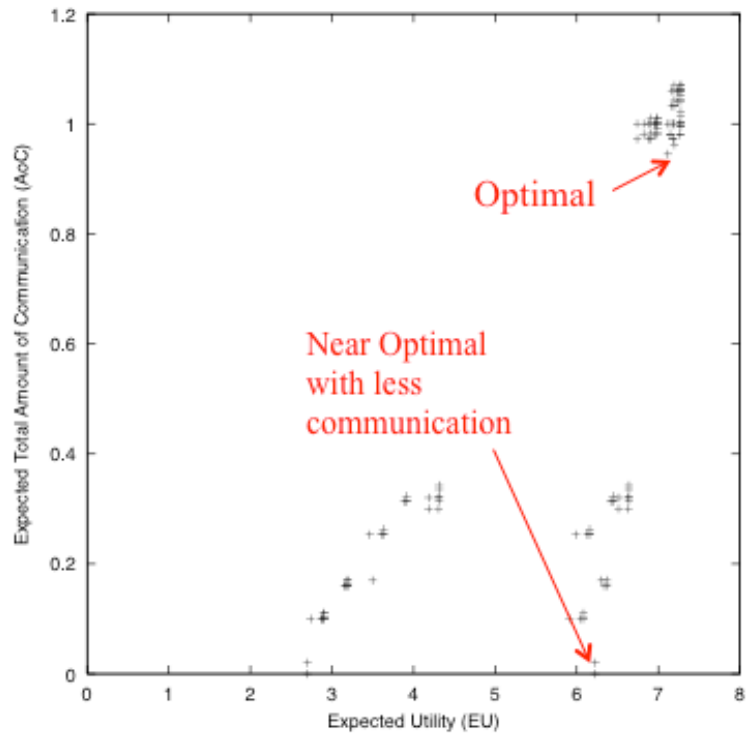


Figure 1: Power of Implicit Communication from [Xuan02:Multi-agent]

Zhang et al. achieved similar results in their more recent work on multi-agent reinforcement learning [Zhang13:Coordinating] where they used a DCOP algorithm to coordinate agent learning to approximate a centralized learning algorithm. In this case, they realized that instead of having one massive DCOP that spanned all agents, they could break the DCOP into a set of much smaller independent DCOPs, which significantly reduced the amount of communication required to implement the coordinated learning, with only a slight reduction in the utility of the learned policies of the agents. In developing this dynamic decomposition of the DCOP, they used a statistical view of agents' states based on their current policy to find situations in which not knowing the current states of specific agents would not significantly decrease other agents' utility—in this way decomposing the network into separable coordination problems. Again, the need for communication among agents did not always relate directly to structural interaction in the problem description, especially when a slight decrease in overall utility was acceptable.

Another case of a distributed problem-solving situation that is more nearly-decomposable than expected (and thus requiring less coordination and communication) was the early work by Lesser

et al. on distributed interpretation/situation-assessment [Lesser80:Distributed, Lesser81:Functionally]. The problem was to construct an overall interpretation of a situation from a group of agents, each having a limited, partial view of it. They developed a successful coordination strategy that created the correct solution a high percentage of the time and that only required a limited exchange of high-level abstract hypotheses generated by each agent. Carver et al. [Carver03:Domain] more rigorously explored why this approach worked and developed the concept of domain monotonicity to explain it. In this case, the power of local problem-solving constraints allowed agents to often be able with strong certainty to narrow down the possible solutions to its part of the overall problem to a small number of cases without knowing the possible solutions to other agents subproblems. This made the amount of communication among agents to find the correct overall interpretation much less than expected. Therefore, the distributed interpretation problem was more loosely connected than what would appear by simply taking into account constraints among agents' subproblems.

We hypothesize that there is something going on that has not been modeled by the explicit structural relationships on agent activities as defined by the problem description. It is not the existence of all interaction relationships that needs to be modeled, but something more nuanced where a trade-off between optimality and communication can be expressed. A theory is needed that connects the characteristics of the problem description and the character of optimal or near-optimal coordination strategies. When only the agents' key interactions (those that can potentially affect overall system utility significantly) and their partial-ordering are considered in the context of likely joint agent states, then agents often are more loosely-connected (more nearly-decomposable [Simon69:Sciences]) than would be expected by the existence of all structural interactions among agents.

Coordinating Agents' Local Computation

Another way of thinking about the observations in the previous section is in terms of what assumptions one agent can make about the state of other agents with whom they potentially interact. In learning theory, this idea is discussed in terms of the concept of a *non-stationary* environment: the more non-stationary the environment is, the harder the learning. Thus, if coordination techniques can decrease or change the nature of the non-stationary environment in multi-agent learning caused by concurrent learning in neighboring agents, then they can improve learning performance significantly in terms of both the speed of convergence and the likelihood that convergence will actually occur. It is our hypothesis that one of the underlying reasons why approaches developed by the multi-agent reinforcement community are effective is that they make the local agent learning algorithm change in slower and more predictable ways [Bowling02:Multiagent, Abdallah08:Multiagent, Zhang10:Multi-Agent]. In this way, even though individual agent learning may not be as efficient from a local perspective, learning from a system-wide perspective can converge more quickly and to better solutions.³

³ The multi-agent learning community has also successfully used the learning of stochastic MDP policies instead of regular MDP policies to speed up convergence. It is our belief that learning these stochastic policies is effective since they act to slow down the rate of policy change that has the effect of damping down the non-stationary character of the environment.

This issue of a non-stationary environment also occurred in different guises in earlier work on developing both heuristic and formal coordination strategies. These examples have an interesting connection with the multi-agent reinforcement learning example discussed above: they all involve the use of iterative algorithms, where the same basic process is repeated on each cycle as new information is received. Brooks et al. [Brooks79:Distributed] early on coined the term “simultaneous-update uncertainty” to describe this non-stationary environment characteristic. They worked on the problem of distributed traffic light control using a distributed iterative algorithm and developed such techniques as modulating the magnitude of changes on any cycle, giving priority to certain neighboring traffic light agents' information changes over other agents' information, and modulating the frequency of updates based on the state of the agents' current traffic control pattern. All of these strategies decreased simultaneous-update uncertainty and improved performance. Similarly, Fernandez, et al., found that the “active introduction of message delays by agents can improve performance and robustness while reducing the overall network load” for distributed constraint satisfaction algorithms (DCSP) [Fernandez02:Communication] (see Figure 2).

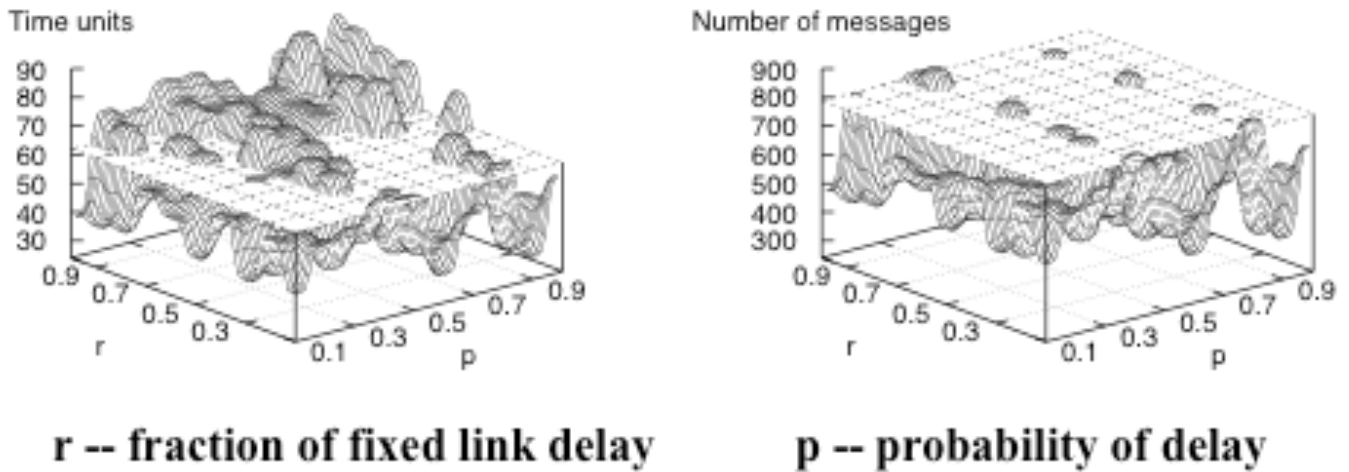


Figure 2: Median time and number of messages to solve hard satisfiable constraint problems when agents add random delays in outgoing messages. The horizontal plane represents when no delay is added ($p=0, r=0$) from [Fernandez02:Communication].

Our hypothesis for explaining this behavior relates to how an iterative improvement search process works. If the search is started with a tentative solution that is partially correct, performance improves significantly. However, even without a good starting point, this type of search can still be effective because it can often find tentative solutions quickly that contain fragments/partial-solutions that correspond to fragments of the correct solution. These correct

fragments direct the search process to find the correct solution; generally, the larger the consistent fragment the quicker the search will progress since a larger fragment often contains more constraints that in turn limit the ways that the fragment partial solution can be extended. For example, consider a distributed search such as asynchronous weak-commitment search (AWC) [Yokoo95:Asynchronous] used by Fernandez, et al. [Fernandez02:Communication], where each agent is solving a component of the overall problem. If the coordination does not allow clusters of agents to form consistent fragments across agents with sufficient frequency due to non-stationarity issues (because agents are frequently switching what they consider as their best current local solutions), then the distributed search will take much longer and, in the case of complete algorithms such as AWC and, in the case of incomplete algorithms, lead to oscillation or convergence to suboptimal solutions. Thus, by slowing down the frequency of updates, it is more likely that groups of agents will construct consistent larger fragments of the overall solution. This will in turn speed up the overall search process. Another way of framing this is that all these heuristic approaches are intended to reduce the oscillation during search caused by concurrent learning or local partial solution update. In some sense, they serialize or coordinate agents' local search activities to improve the performance [Zhang13:Private].

This behavior is part of a larger phenomenon, called “distracting communication” [Lesser80:Distributed], which involves agents transmitting information that is no longer appropriate for or relevant to the receiving agent since it is in a different state of problem-solving, or because the information is not correct or out-of-date or the likelihood measure associated with its correctness is not accurate. These distracting communications, if they occur with sufficient frequency, then have the effect of continually diverting overall system problem-solving to unprofitable paths that lead to longer problem-solving times due to more backtracking of problem-solving, more communication, and less accurate or no solutions. Lesser et al. [Lesser80:Distributed] suggested a number of heuristics for handling this problem such as: 1) delaying sending information until more problem-solving has been completed so that it is easier to assess the validity of the information, 2) having better local problem-solving strategies that reduce the likelihood of sending incorrect information, and 3) exchanging meta-information among agents to better gauge the current problem-solving directions of other agents.

Unfortunately, a formal and quantitative theory of distributed search that explains in detail why the above heuristic approaches to improving agent coordination work has yet to be developed.⁴ What is missing is a theory that explains how both the character and frequency of incorrect or out-of-date information affects the performance of a coordination strategy and, ultimately, overall network problem solving.⁵ The difficulty in developing such a formal theory is that the consequences of this inaccurate information and associated problem solving are not confined to individual agents but can propagate throughout the agent network. Thus, there is a need to incorporate some type of statistical model of the distributed search process being coordinated and

⁴ There is limited formal analysis showing why such approaches work in multi-agent reinforcement learning community [Bowling02:Multiagen, Abdallah08:Multiagent, Zhang10:Multi-Agent] where there are only two agents interacting.

⁵ Zhang et al. [Zhang10:Multi-Agent] present a formal approach for talking about interactions among agent policies that takes into account frequency and strength of interactions; however, they did not deal with how interactions among agents propagate in the network nor the specifics of the coordination strategy.

its associated intermediate states into a formal analysis framework for explaining these phenomena⁶.

Coordination and Environment

The issue of moderating the frequency of coordination has also come in another guise. Durfee et al. [Durfee88:predictability] introduced the trade-off between predictability and responsiveness, where communication and computation costs associated with coordination are modulated by varying the conditions when an agent indicates that its current state does not match the expectations used in the current coordination strategy. In this case, a wider tolerance for variance from the expected agent behavior leads to more predictability in a coordination strategy (since it is less likely to be revised) with the consequence that the strategy is not as responsive to the details of the current agents' states and thus the coordination is not as precise. However, they observed that, given the additional costs and delays of being highly responsive, it may be better to use a less responsive coordination strategy.⁷ This example of what Simon [Simon69:Sciences] called "satisficing" in which optimal decision making is not always preferable over near-optimal decision making given the associated costs (in this case communication and computation, and associated delays) of making such optimal decisions. We would argue that this satisficing approach works for multi-agent systems because most distributed application agent activities are more loosely-connected than would be expected based on structural interactions and therefore incorrect decisions that do not take into account all structural interactions are often not catastrophic and can be either corrected downstream or do not affect overall performance significantly.

More generally, depending on the characteristics of the environmental conditions in terms of resource availability, task loading, and predictability of task behaviors, very different coordination strategies are appropriate. Without going into detail, here are our summaries of some of the observations. The first observation is that in environments with very high or very low task loading or high variance in agent behavior, simple coordination strategies work quite well.⁸ However, this does not contradict our basic point because the specific instances of high variance in this case could be ascertained before the coordination strategy was constructed (based on meta-level information) rather than needing to be recognized during the execution of the coordination strategy. It is only in situations where there is a "sufficient" level of predictability about agent behavior or intermediate levels of task loading that complex coordination strategies are advantageous. Corkill et al [Corkill15:Exploring] call this the "sweet spot". This last point relates to the nature of phase transition, where the difficulty of solving problems increases significantly around the phase transition, and effective coordination can make the difference in satisfactory

⁶ We are familiar with some work that has a formal character that can predict overall system performance and is able to describe heuristic control knowledge and how it affects the underlying search strategy, but this formal work is for a specific search strategy and only for a single agent [Whitehair95:Thesis].

⁷ Further, they found in one case that, even if these additional costs were discounted, it remained better to be more predictable because the coordination strategy was constantly changing on each cycle, causing unnecessary backtracking of agent problem solving in a way similar to what we discussed.

⁸ Decker et al. showed formally, for a specific task allocation problem, that if there was high variance in the number of tasks associated with different agents more sophisticated coordination strategies that exploited meta-level control information did better [Decker93:Approach].

performance. We suggest that there are similar phase transitions going on in agent coordination and that it is in those transition regions where more complex control is advantageous (see Figure 3).



Figure 3: The Effects of Organizational Control in Different Task Environments from [Corkill15:Exploring].

The second observation is that even though a Dec-MDP can be used to build an optimal coordination strategy, simpler non-optimal heuristic coordination strategies that only consider the major interactions among agents and do not deal with contingencies directly (or deal with them in only limited ways) do quite well in most coordinating situations (see [Lesser04:Evolution, Wagner03:Key-based]). These non-optimal approaches re-coordinate when necessary based on the actual contingent event rather than attempting to prevent such situations from occurring or planning ahead for all contingencies. The hypothesis behind this observation is again that agent interactions in most situations are more loosely connected than would be expected, and most incorrect coordination decisions can be tolerated and corrected without severe harm to overall agent performance. The question for us is whether there is a formal way for looking at a problem and its environmental description to understand what contingencies, agent activity horizons, and problem-solving states of other agents need to be considered in order for coordination to work effectively.

We have discussed above that, depending on the environmental characteristics, very different approaches to coordination are appropriate. The same holds for the underlying distributed search or learning process being coordinated. In this case, the character of the distributed search process for a single problem may vary significantly over its lifetime. Mailler et al. recognized this in developing a very effective approach to distributed constraint satisfaction where the scope of control (partial centralization of control) varied on each cycle based on the current constraint interactions among the partial solutions constructed at different agents [Mailler06:Asynchronous]. In this case, the partial centralization was introduced to handle situations where the solution to a subproblem associated with subset of agent constraints required those agents to change their current local solutions in a way that violated the constraints of other interconnected agents that were not directly involved in the subproblem being solved; this situation then led to a new partial centralization of control that considered the constraints associated with these other agents as part of a new subproblem that needed to be solved. This approach decreased the likelihood of backtracking that occurs in normal DCOP search. Similarly, Zhang et al. used a strategy for coordinating multi-agent learners that dynamically change their scope of control based on the strength of interaction among current learned policies of agents [Zhang13:Coordinating, Zhang10:Self-organization]. More generally, a coordination strategy that can adapt to the current situation seems crucial where the environment or network problem solving is evolving dynamically and rapidly, and different situations require different approaches to coordination.

Conclusions

Our intuition is that all these experimental behaviors and phenomena are interrelated and that an integrated and formal treatment of them by future generations of researchers will lead to a much deeper understanding of the nature of coordination and cooperation and, more generally, decentralized control. Lacking from current formal frameworks are: 1) a statistical model of the underlying distributed search process (network problem solving) that is being coordinated and its associated intermediate states and 2) formal treatment of concepts such as “nearly-decomposable” and “satisfiability” developed by Simon [Simon69:Sciences] for understanding the relationship between effective coordination and acceptable but non-optimal performance. That is our challenge to the multi-agent field.

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under Award Numbers IIS-0964590 and IIS-1116078. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [Abdallah08:Multiagent] S. Abdallah and V. Lesser. A Multiagent Reinforcement Learning Algorithm with Non-linear Dynamics. *Journal of Artificial Intelligence Research*. Volume 33, pages 521–549, 2008.
- [Becker04:Solving] R. Becker, et al. Solving transition independent decentralized Markov decision processes. *Journal of Artificial Intelligence Research*, 22:423–455, Jul–Dec 2004.
- [Bernstein00:Complexity] D. S. Bernstein, S. Zilberstein, and N. Immerman. The complexity of decentralized control of Markov decision processes. In *Proceedings of the 16th International Conference on Uncertainty in Artificial Intelligence*, pages 32–37, Stanford, California, July 2000.
- [Bowling02:Multiagen] M. Bowling and M. Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136:215–250, Apr. 2002.
- [Brooks79:Distributed] R. S. Brooks and V. R. Lesser. Distributed problem solving using iterative refinement. Technical Report 79-14, Department of Computer and Information Science, University of Massachusetts Amherst, Amherst, Massachusetts 01003, May 1979.
- [Davis83:Negotiation] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, pages 63–109, 1983.
- [Decker93:Approach] K. Decker and V. Lesser. An approach to analyzing the need for meta-level communication. In *IJCAI 1993*, pages 360–366, Chambéry, France, Aug. 1993.
- [Durfee88predictability] E. Durfee and V. R. Lesser. Predictability versus responsiveness: Coordinating problem solvers in dynamic domains. In *AAAI 1988*, pages 66–71, St. Paul, Minnesota, Aug. 1988.
- [Durfee91:Partial] E. H. Durfee and V. R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-21(5):1167–1183, May 1991.
- [Carver03:Domain] Carver, N. and Lesser, V.R. Domain Monotonicity and the Performance of Local Solution Strategies for CDPS-based Distributed Sensor Interpretation and Distributed Diagnosis. *Autonomous Agents and Multi-Agent Systems*, Kluwer Academic Publishers, 6(1): 35–76, 2003.
- [Corkill15:Exploring] D. Corkill, D. Garant and V. Lesser. Exploring the Effectiveness of Agent Organizations. *Proceedings of COIN Workshop, AAMAS'15*, May 2015.
- [Farinelli08:Decentralized] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings. “Decentralised coordination of low-power embedded devices using the max-sum algorithm” In *AAMAS*, 2008, pages 639-646. 2008.
- [Fernandez02:Communication] Cesar Fernandez, et al. Communication and computation in distributed csp algorithms. In V. Lesser, C. L. Ortiz, Jr., and M. Tambe, editors, *Distributed Sensor Networks: A Multiagent Perspective*, chapter 12, pages 299–318. Kluwer Academic Publishers, 2003.

- [Goldman04:Decentralized] C. V. Goldman and S. Zilberstein. Decentralized control of cooperative systems: Categorization and complexity analysis. *Journal of Artificial Intelligence Research*, 22:143–174, 2004.
- [Jennings93:Commitments] N. R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250, 1993.
- [Kumar11:Scalable] A. Kumar, S. Zilberstein, and M. Toussaint. Scalable multiagent planning using probabilistic inference. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2140–2146, 2011.
- [Lesser80:Distributed] V. R. Lesser and L.D. Erman. “Distributed Interpretation: A Model; an Experiment.” *IEEE Transactions on Computers - Special Issue on Distributed Processing*, Vol. C-29, 12, pages 1144–1163, 1980.
- [Lesser81:Functionally] V. R. Lesser and D. D. Corkill. Functionally accurate, cooperative distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1):81–96, Jan. 1981.
- [Lesser91:Retrospective] V. Lesser. A Retrospective View of FA/C Distributed Problem Solving. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6): 1347–1362, 1991.
- [Lesser98:Reflections] V. R. Lesser. Reflections on the nature of multi-agent coordination and its implications for an agent architecture. *Autonomous Agents and Multi-Agent Systems*, 1(1):89–111, March 1998.
- [Lesser04:Evolution] V. R. Lesser, et al. Evolution of the GPGP/TÆMS domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems*, 9(1):87–143, July 2004.
- [Macarthur10:Superstabilizing] K. Macarthur, A. Farinelli, S. Ramchurn and N. Jennings. Efficient, Superstabilizing Decentralised Optimisation for Dynamic Task Allocation Environments. *Proceedings of the 20th National Conference on Artificial Intelligence*, pages 449–454, 2010.
- [Mailler06:Asynchronous] R. Mailler and V. R. Lesser. Asynchronous partial overlay: A new algorithm for solving distributed constraint satisfaction problems. *Journal of Artificial Intelligence Research*, 25:529–576, Jan–Apr 2006.
- [Mostafa11:Compact] H. Mostafa and V. Lesser. “Compact Mathematical Programs For DEC-MDPs With Structured Agent Interactions.” *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*, Barcelona, pages 523–530, 2011.
- [Mostafa11:Private] H. Mostafa. Private communications on the difficulty of finding optimal policies that benefit from explicit communication for the Mars Rover scenarios described in her Ph.D. dissertation, 2011.
- [Monasson99:Determining] R´emi Monasson, et al. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400:133–137, July 1999.
- [Nair05:NETPOMDPS] R. Nair, P. Varakantham, M. Tambe, and M. Yokoo. Networked distributed POMDPs: A synthesis of distributed constraint optimization and POMDPs. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, pages 133–139, 2005.

- [Pujol2013:Binary] M. Pujol-Gonzalez, J. Cerquides, G. Escalada-Imaz, P. Meseguer, J. Rodriguez-Aguilar. [On Binary Max-Sum and Tractable HOPs](#), 11th European Workshop on Multi-agent Systems (EUMAS 2013), Volume 1113, Toulouse, France (2013).
- [Sen1998:Meeting] S. Sen and E. H. Durfee. A formal study of distributed meeting scheduling. *Group Decision and Negotiation*, 7(3):265–289, May 1998.
- [Seuken08:FormalModels] S. Seuken and S. Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *Autonomous Agents and Multi-Agent Systems*, 17(2): 190–250, 2008.
- [Simon69:Sciences] H. A. Simon. *The Sciences of the Artificial*. MIT Press, 1969.
- [Tambe97:Towards] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, Jul–Dec 1997.
- [Tarlow10:HOP] D. Tarlow, I.E. Givoni and R.S Zemel. HOP-MAP: Efficient Message Passing with High Order Potentials. In: 13th International Conference on Artificial Intelligence and Statistics (AISTATS). Volume 9. 812-819, 2010.
- [Wagner03:Key-based] T. Wagner, V. Guralnik, and J. Phelps. A key-based coordination algorithm for dynamic readiness and repair service coordination. In *AAMAS 2003*, pages 757–764, Melbourne, Australia, July 2003.
- [35Whitehair95:Thesis] R. Whitehair. [A Framework for the Analysis of Sophisticated Control](#). Department of Computer Science, *University of Massachusetts – Amherst, PhD Thesis*. 1995.
- [Witwicki11:Towards] S. J. Witwicki and E. H. Durfee. Towards a Unifying Characterization for Quantifying Weak Coupling in Dec-POMDPs. In *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*, pages 29-36, 2011.
- [Xuan02:Multi-agent] P. Xuan and V. Lesser. Multi-agent policies: From centralized ones to decentralized ones. In *AAMAS 2002*, Bologna, Italy, July 2002.
- [23Yeoh13:Automated] W. Yeoh, A.t Kumar, and S. Zilberstein. Automated Generation of Interaction Graphs for Value-Factored Dec-POMDPs. *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, Beijing, China, 2013.
- [Yokoo95:Asynchronous] M. Yokoo. Asynchronous weak-commitment search for solving distributed constraint satisfaction problems. In *Proceedings of the First International Conference on Principles and Practice of Constraint Programming (CP '95)*, pages 88–102, Cassis, France, Sept. 1995.
- [Yokoo98:Distributed] M. Yokoo and E. H. Durfee. Distributed constraint optimization as a formal model of partially adversarial cooperation. Technical Report CSE-TR-101-91, University of Michigan, 1991.
- [Zhang10:Multi-Agent] C. Zhang and V. Lesser. Multi-agent learning with policy prediction. In *AAAI 2010*, pages 927–934, Atlanta, Georgia, July 2010.
- [Zhang10:Self-organization] . Zhang, V. Lesser, and S. Abdallah. Self-organization for coordinating decentralized reinforcement learning. In *AAMAS 2010*, pages 739–746, Toronto, Canada, May 2010.

- [Zhang13:Coordinating] C. Zhang and V. Lesser. Coordinating multi-agent reinforcement learning with limited communication. In AAMAS 2013, pages 1101–1108, St. Paul, Minnesota, May 2013.
- [Zhang13:Private] C. Zhang. Private communication, 2013.
- [Zivan10:Distributed] R. Zivan, R. Glinton and K. Sycara. Distributed Constraint Optimization for Large Teams of Mobile Sensing Agents. 2010. In IAT 2009, pages 347-354.