A Unified Framework for Multi-Agent Agreement

Kiran Lakkaraju
Department of Computer Science
University of Illinois, Urbana-Champaign
klakkara@uiuc.edu

Les Gasser
Graduate School of Library and Information Science
University of Illinois, Urbana-Champaign
gasser@uiuc.edu

ABSTRACT

Multi-Agent Agreement Problems (MAP) - the ability of a population of agents to search out and converge on a common state - are central issues in many multi-agent settings, from distributed sensor networks, to meeting scheduling, to development of norms, conventions, and language. While much work has been done on particular agreement problems no unifying framework exists for comparing MAPs that vary in, e.g., strategy space complexity, inter-agent accessibility, and solution type, and understanding their relative complexities. We present such a unification, the Distributed Optimal Agreement (DOA) framework, and show how it captures a wide variety of agreement problems. To demonstrate DOA and its power we apply it to convention evolution. [1]

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*multi agent systems*

General Terms

Language Evolution, multi agent agreement problems, multi agent systems.

Keywords

Language Evolution, multi agent agreement problems, multi agent systems

1. INTRODUCTION

In a *Multi-Agent Agreement Problem* (MAP) multiple agents must navigate a space of possible states (a *potential agreement space*) and eventually converge on the same state. By illustration, a simple MAP from distributed transaction processing is the *distributed commit* problem: all agents participating in a single transaction must eventually converge on one state in a two-valued potential agreement space, namely whether to *commit* or *abort* the transaction. The well-known "two-phase commit" and "three-phase commit" protocols are different specific *solutions* to this MAP, with differing degrees of robustness, complexity, and centralization (cf. [4]).

Different types of MAPs occur in many other multi-agent domains as well, such as distributed active sensor networks. Similar problems arise in stabilizing the physical formation of a set of agents (agreement on roles and positions), flocking/swarming (agreement on overall direction), and synchronization of coupled oscillators (agreement on state of the oscillation). Negotiation can be viewed as a MAP, in which agents aim to converge on a common state in an *offer space* (or fail). Meeting scheduling can be viewed as a MAP governed by a set of equality constraints over time and availability. Multi-agent agreement is at the heart of the notions of "norms" and "conventions" for multi-agent systems [7]. Here the potential agreement space is the collection of agents' possible action strategies in a particular class of situations, and the resulting convention is a strategy agreed to and followed by all agents.

Our particular area of interest is the autonomous creation of language by an agent collection [1, 2, 8, 9][2]. Language is necessarily a collective phenomenon. One quality measure over a population of language users is the population's average *communicability*[3]. Communicability implies agreement on language. A set of agents can achieve high communicability by settling on any one of a wide variety of possible languages. But individual languages might also differ on "objective" qualities, such as *expressivity* (ability to express important concepts), *efficiency* (of production and interpretation), and so on. Thus, for the language MAP, a solution has both a *frequency-dependent* component (communicability, which depends on the *frequency* of agents sharing a language) and an *intrinsic* or objective quality component. Agents need *agreement on a common and high-quality state (language)*.

At first glance it would seem that since these various problems can all be cast as agreement problems, both the MAPs and their solutions should follow a general and by now well understood pattern. Indeed, our first assumption was that

[1]For a more detailed description of this work see [3]

[2]Language is also important because it is a model problem for dynamic semantics in general in distributed information systems: e.g., evolving shared schemas/ontologies, shared denotations, common information structures, etc.

[3]Defined, e.g., as the average probability that an "utterance" by one agent will be understood by another; see [6]).

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

well-known distributed algorithms such as those collected in [4] provide the solution concept for most MAPs, so we could reuse these and little more remained to be done. Trying this, however, we found that each MAP domain has its own idiosyncrasies of problem description that are sometimes hard to apply to other domains. Some kind of MAP *lingua franca* is needed to validate any unity among MAPs. Our contribution in this paper is to present a framework, which we call *Distributed Optimal Agreement (DOA)*, that can unify the different agreement problems from different problem domains under one common descriptive model, allowing researchers to compare, contrast, and understand MAP differences clearly. We detail the DOA framework in Section 2 and show an example of a MAP captured with the Distributed Optimal Agreement framework in Section 3

2. DOA FRAMEWORK

The Distributed Optimal Agreement framework comprises a specification of a problem and a specification of the dynamics of the agents in solving that problem. Informally, we view the process of solving a MAP as collective search. Each agent moves about in a *possible agreement space* that comprises a number of *possible agreement states (PAS)*. Any of the PASes might be the substance of an agreement, depending on its own qualities and the number of agents that have settled on it. A *complete agreement* is the condition that all searching agents have arrived at the same PAS. DOA problems can differ on many characteristics, including *complexity of the potential-agreement space*, *state-to-state accessibility*, *agent interaction topologies*, *state evaluation measures*, and *type of solution needed* for instance.

We present the formal DOA model below.

2.1 Formal Problem Model

An agreement problem in the DOA framework is defined by the 7-tuple:

$$\{A, \Sigma, \Delta, \Theta, \rho, S, \Omega\}$$

where:

1. **Agents:** A is a set of n *agents*, $\alpha \in A$. Agents are the active processes in the DOA model, whose actions take place in an interval in a time line T. At each time $t \in T$, an agent is said to be "in" some Possible Agreement State (see below).

2. **Possible Agreement Space:** The substance of an agreement in the DOA model is the *possible agreement state (PAS)*, denoted by σ: a state of the world on which agents could agree. For instance, a PAS could be a language an agent chooses to speak, an offer in a negotiation, a candidate strategy for a convention, or a decision to commit/abort a transaction. Some agreement spaces such as team coordination can be very complex, and hence difficult to make into metric search spaces; they are nonetheless search spaces and modelable with DOA.

   Σ is the set of all PASs, thus $\sigma \in \Sigma$. We use $\sigma_{\alpha_i,t}$ to denote the PAS that agent α_i is "in" at time t.

   Configurations

   Let Σ^n be the set of all possible associations of PASs with all the agents in A. Σ^n is thus an n-dimensional space. At time t the *configuration* of the entire system is $s_t \in \Sigma^n$—that is, one specific association of all agents with states.

3. **Accessibility Relation:** $\Delta : A \times \Sigma \times \Sigma \rightarrow \{\Re \cup \infty\}$ is the *accessibility relation* for PASes σ. $\Delta(\alpha_i, \sigma_j, \sigma_k)$ describes the (possibly infinite) cost for some agent α_i to move from σ_j to σ_k. Δ models the structure of the possible agreement space Σ from the perspective of each agent. An agent with more limited capabilities might have a higher cost for changing from one PAS to another, or one PAS might be inherently more difficult (or impossible) to reach directly.

4. **Interaction Relation:** $\Theta : A \times A \times T \rightarrow \{\Re \cup \infty\}$ is the *interaction relation*. $\Theta(\alpha_i, \alpha_j, t_i)$ describes the cost for an agent α_i to interact with (e.g. sense, observe, communicate with) some other agent α_j at time $t_i \in T$. Cost is a very general basis for an interaction relation. For example, a close interaction neighborhood for some agent can be defined as the set of other agents with which communication is cheap relative to other agents. If cost is inversely related to probability of interaction over time, then Θ describes agent-to-agent interaction frequencies, and can be used to model a type of *frequency-weighted social network*.

5. **Intrinsic Value:** $\rho : A \times \Sigma \rightarrow \Re$ defines the *intrinsic value* of an agent being in a particular PAS. $\rho(\alpha_i, \sigma_j)$ defines the reward agent α_i receives from being in PAS σ_j. Σ can be seen as a landscape with hills and valleys corresponding to ρ. Since ρ is defined based only on the agent and what strategy it is using, and not on what strategies other agents have, we consider ρ to be the *intrinsic* value of the state with respect to an agent. In many cases ρ is independent of the agent as well. We define $max(\rho)$ as the set $\{(\alpha_i, \sigma_j)\}$ with the highest $\rho(\alpha_i, \sigma_j)$

6. **Starting Configurations** The set of possible starting configurations, $S \subseteq \Sigma^n$. $s_0 \in S$ is the initial state of the population.

7. **Termination Configurations** The set of possible termination configurations, $\Omega \subseteq \Sigma^n$. There are many types of termination configurations. Here are several interesting ones:

   Simple Consensus Configurations in Ω are *agreements*. A *complete agreement* is formed by a set of agents all being "in" the same PAS, for example all choosing to subscribe to a *particular* language, negotiation offer, convention strategy, etc. This is denoted as a configuration with the following property:

   $$s \ni \forall i, j, \sigma_{\alpha_i,t} = \sigma_{\alpha_j,t} \tag{1}$$

   Consensus+Optimization At some t $\sigma_{\alpha_i,t} = \sigma_{\alpha_j,t}$, $\forall i, j$ and $(\alpha_0, \sigma_{\alpha_0,t}) \in max(\rho)$. This is the set of configurations in which every agent is using the same strategy, and that strategy has the highest intrinsic value.

   Consensus+Computation Given a function $\chi : \Sigma^n \rightarrow \Sigma$, at some t, $\sigma_{\alpha_i,t} = \chi(s_0)$, $\forall$ i. The set of configurations where every agent is in the same PAS, and that specific PAS is a function of the initial PASes of the entire population.

2.2 System Dynamics

Solving an instance of a DOA problem involves specifying the behavior of the agents such that the system moves from a configuration $s_0 \in S$ to a configuration $s_\omega \in \Omega$ in some finite amount of time. We assume a turn-based system, where at each time step t the three-step process of active agent selection, information gathering, and information use occurs, as follows:

1. Active Agent Selection: A subset $C_t \subset A$ (called the *active agent set*) becomes active at this time step.

2. Information Gathering: Information (call it ψ) is necessary for efficient search. Complete information about a system configuration is costly, being influenced by Θ, n, and $|\Sigma|$ and the history of activity represented by T. Thus strategic selection of a subset of information sources at each time step is necessary in a MAP. Once this choice is made, the actual interactions occur.

2a. Interaction Choice Each active agent $\alpha_i \in C_t$ chooses some other subset of agents $I_{i,t} \subset A$ (called the *interaction set* of α_i), from which to gather information about the current configuration. This substep is purely the choice of a set of other agents to observe or communicate with.

2b. Interaction) α_i interacts with the agent(s) in I from the previous step. This interaction produces some *information* for α_i about the current configuration. We note here that our empirical study of a wide variety of MAPs and solution concepts showed how the *completeness of an agent's knowledge of other agents' states* is a critical element distinguishing the effectiveness of different MAP solution concepts.

3. Information Use Finally, all agents active in this time step (C_t) apply a decision rule $\sigma_{\alpha_i,t+1} = f(\sigma_{\alpha_i,t}, \psi)$ possibly moving to another PAS.

3. CONVENTION EMERGENCE IN THE DOA FRAMEWORK

The DOA framework is general enough to represent many different situations described in the MAS community. As an example, we show how to map convention emergence situations (as described in [7]) into the DOA framework.

Shoham & Tennenholtz (ST) [7] studied the emergence of *social conventions* - social laws that restrict agents' behavior and so enforce a particular global strategy. A social convention must be agreed upon by every agent in the society, thus it is an instance of a multi-agent agreement problem. In [7] ST used a *stochastic game* framework to explore the emergence of social conventions. Theirs was a turn based system where at each time step two agents were chosen from the population to interact in a 2-person-2-choice symmetric game: the *coordination game*, in which agents received a positive payoff if both agents played the same strategy, and negative payoff otherwise. [7] showed that social conventions emerged in the coordination game when agents used a *Highest Cumulative Reward* (HCR) rule to choose their strategy. HCR causes an agent to switch to the one strategy (out of just two choices) that has provided the maximum cumulative reward over the past window of m interactions.

Mapping the Shoham & Tennenholtz model into the DOA framework is straightforward:

Agents Population of agents.

Agreement Space $\Sigma = \{0, 1\}$ - the space of strategies an agent can play.

Accessibility Relation $\Delta(\sigma_i, \sigma_j) = 0\ \forall\ \sigma_i, \sigma_j$ (no restrictions).

Interaction Cost $\Theta(\alpha_i, \alpha_j) = 0\ \forall \alpha_i, \alpha_j$ (no restrictions).

Intrinsic Value $\rho(\alpha_i, \sigma) = 0\ \forall \alpha_i, \sigma$ (all strategies equivalent).

Start Configurations Agents start in random states.

Termination Configuration All agents must be in the same state - the Consensus condition.

Modeling their solution is also straightforward. At each time step two agents are picked to be active. This is done uniformly randomly with no agent influence. Both C_t and I_t are chosen in one step. In the information gathering stage the two agents play a stochastic game defined by a payoff matrix. The payoffs do not indicate the state of the other agent, so we can consider this a case of gathering only *indirect* knowledge of the other agents state. Finally, the agents employ the HCR rule to update their state. This is mapped into the information use stage of the system dynamics.

4. CONCLUSION

We have developed the Distributed Optimal Agreement framework to describe a wide variety of multi-agent agreement problems in order to compare, contrast and understand the differences in complexity of MAP problems. We have shown how a particular MAP problem (the emergence of social conventions) can be mapped into the DOA framework.

5. REFERENCES

[1] F. Cucker, S. Smale, and D.-X. Zhou. Modeling language evolution. *Foundations of Computational Mathematics*, 4(3):315–343, July 2004.

[2] K. Lakkaraju and L. Gasser. The complexity of finding an optimal policy for language convergence. In *Proceedings of the the Ninth International Conference on the SIMULATION OF ADAPTIVE BEHAVIOR (SAB'06)*, 2006.

[3] K. Lakkaraju and L. Gasser. A unified framework for multi-agent agreement. Technical Report UILU-ENG-2006-1843, University of Illinois, Urbana-Champaign, 2007.

[4] N. A. Lynch. *Distributed Algorithms.* Morgan Kaufmann, 1997.

[5] F. A. Matsen and M. A. Nowak. Win-stay, lose-shift in language learning from peers. *PNAS*, 101(52):18053–18057, December 2004.

[6] M. A. Nowak, N. L. Komarova, and P. Niyogi. Computational and evolutionary aspects of language. *Nature*, 417(6889):611–617, June 2002.

[7] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1-2):139–166, July 1997.

[8] L. Steels. The origins of ontologies and communication conventions in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 1(2):169–194, October 1998.

[9] K. Wagner, J. Reggia, J. Uriagereka, and G. Wilkinson. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11:37–69, 2003.