ETAPP: A Collaboration Framework That Copes with Uncertainty Regarding Team Members*

Christian Guttmann

School of Computer Science and Software Engineering,
Monash University, Clayton, VICTORIA 3800, Australia
xtg@csse.monash.edu.au

1 Introduction

The organized nature of human collaboration is often used as a metaphor for computational theories of collaboration. Knowledge of collaborators' capabilities and reliability of decision making processes are important factors in collaborative activities. In this thesis, we investigate these factors in the context of an important collaborative activity – the assignment of team members to tasks.

Many collaboration theories assume that agents have correct and complete knowledge of the capabilities of team members as well as team members that make decisions in a reliable manner. Making decisions in a reliable manner means that decisions are made to optimize utility according to the criteria of a task (rather than an agent's own criteria). In the assignment of agents to activities, team performance is optimized by assigning optimal agents (whose capabilities are known) to activities.

Consider a simplified scenario where a doctor refers patients to specialists. Assume the doctor knows the specialists' capabilities and makes decisions to optimize the treatment of patients. Based on these assumptions, the doctor would refer patients to optimal specialists in order to obtain optimal treatment. However, we argue that such referrals are not optimal when a doctor is not fully aware of the capabilities of specialists and/or makes decisions according to the doctor's own criteria. Consider some examples in support of our argument.

- **Limited reasoning capabilities** of decision makers. A doctor can not take into account all specialists when making a referral, because a doctor can only remember and assess the capabilities of a limited number of specialists.
- **Variable performance** of specialists. A doctor can make a wrong recommendation based on the overall performance of a specialist, because a specific performance is due to factors that are not known to the doctor (e.g., a specialist has a bad day).
- **Decision making** in an unreliable manner. A doctor can make bad referrals, because the doctor is lazy, meaning that the doctor does not deliberate about which specialist could be optimal, thus selecting an arbitrary specialist.

* This research was supported in part by Linkage Grant LP0347470 from the Australian Research Council, and by an endowment from Hewlett Packard. I would like to thank my supervisor Ingrid Zukerman, and also Michael Georgeff for his advise on this paper.

L. Ardissono, P. Brna, and A. Mitrovic (Eds.): UM 2005, LNAI 3538, pp. 518–521, 2005.

The assumptions of correct and complete knowledge of collaborators' capabilities and reliability of decision makers simplify the problem of agent collaboration. Relaxing these assumptions raises several issues.

- How should agent performance be modelled? How do the limited reasoning capabilities of agents influence model accuracy? How do model accuracy and variable agent performance influence collaborative activities?
- What constitutes decisions that are made in an unreliable manner? How does the unreliability of decision makers influence collaborative activities?
- A team may decide to make group decisions to cope with agents that have incomplete knowledge and agents that make decisions in an unreliable manner. Which kind of group decision procedures improve the quality of collaborative activities? What are the transaction costs involved to arrive at a group decision? What are the tradeoffs between transaction costs and team performance? How many reliable team members are required for robust performance?

These questions are addressed as part of the research in this thesis. Our *research goal* is to analyze collaborative multi-agent behaviour, determine the factors that influence team performance, make predictions of the outcome of team performance, and offer guidelines for efficient collaboration. To this end, we have designed a framework according to a view of agent collaboration which focuses on modelling agents and coping with decisions that are made in an unreliable manner.

Apart from medical referral scenarios, our approach can be applied to a number of examples. For example, making allocations in peer-to-peer networks, where a group of peers is selected based on their bandwidth to establish a transmission routing between two remote peers. Single peers have incomplete knowledge of the current bandwidth of peers in the network, because peers may have insufficient capabilities to monitor every peer, and peers may have insufficient memory to store information of all peers in a network. Additionally, single peers can make decisions that do not optimize the overall performance of a network, instead they make decisions that saves their own bandwidth.

2 Related Research

The development of our framework involved research on a number of topics – agent modelling/tracking/monitoring, decision making, distributed control, and social choice. Our research project is different to other research projects in terms of the research goal, which defines assumptions and agent features considered by an approach. For example, according to Suryadi and Gmytrasiewicz [4] agents do not communicate, thus relying on observations to model other agents. This framework contrasts with research where agent models are derived only from communication [1, 5]. In each project, agents use models of different agent features. For example, expertise is used in student support environments [5], and availability is important in agenda scheduling [1].

3 ETAPP: A Framework of Agent Collaboration

We define a framework called *ETAPP (Environment-Task-Agents-Policy-Protocol)*, that offers an approach to cope with large numbers of collaborating individuals [2, 3, 6]. This framework expresses the collaboration of a team of agents in terms of five operating parameters: Environment, Task, Agents, Policy and Protocol. Briefly, the *Task* given to a group is to be performed in the *Environment*, and the *Policy* and *Protocol* are procedures agreed upon by all the agents in a group, but performed autonomously (this is similar to abiding by certain rules in order to belong to a society). Central to the ETAPP framework is the idea that the team members do not know the real capabilities of the agents in a team. Hence, individual agents employ models of collaborators' capabilities in order to estimate the value of contributions of team members to a task. We also examine agents that make decisions in an unreliable manner, meaning that agents can make decisions that do not optimize utility according to the criteria of a task (but according to an agent's own criteria). The *Agents* component describes a group of agent where each agent stores these models and uses mechanisms to reason about them.

4 Current Contributions and Thesis Schedule

Our specific contributions to research in agent collaboration include the following.

- **A collaboration framework** which offers novel contributions.
  - A probabilistic representation of agent performance in terms of the evaluation criteria of a task (such as time or quality).
  - Cost functions that measure the extent to which communication, computation and memory is used in each collaboration.
  - Voting policies that aggregate decisions of individual agents.
- **Insights, predictions and guidelines**. Based on empirical studies, we found that
  - Several reasoning limitations influence the performance of a team and the transaction costs of a collaboration, e.g., memory and the ability to learn are the most influential factors of team performance and transaction costs [2].
  - Appropriate policies should be used to cope with selfish, conservative, lazy, and corrupt agents, e.g., if agents make decisions in a reliable manner, group decision policies should be simple to improve the performance of the team [3].
  - Variability of individual agent performance influences team performance, e.g., the more variable the agent performance the worse the team performance [6].

We propose the following extensions to our research.

- **Investigate models of team performance** as an extension of our current implementation, where we consider models of only one agent.
- **Decentralize the evaluation of performance**. This means that each agent uses a different function to evaluate observed performance, as opposed to our current approach where one evaluation function is used by all agents.
- **Compare central with distributed decision making procedures**, specifically decisions made by a leader, or decisions derived from voting and auctioning.
- **Provide guidelines** on balancing transaction costs against task performance.

References

1. Leonardo Garrido, Katia Sycara, and Ramon Brena. Quantifying the Utility of Building Agents Models: An Experimental Study. In *Proceedings of the Agents-00/ECML-00 Workshop on Learning Agents*, Barcelona, Spain, 2000.
2. Christian Guttmann and Ingrid Zukerman. Towards Models of Incomplete and Uncertain Knowledge of Collaborators' Internal Resources. In Jörg Denzinger, Gabriela Lindemann, Ingo J. Timm, and Rainer Unland, editors, *Second German Conference on MultiAgent system TEchnologieS (MATES) 2004*, LNAI 3187, 58–72, Erfurt, Germany, 2004. Springer.
3. Christian Guttmann and Ingrid Zukerman. Voting policies that Cope with Unreliable Agents. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, Utrecht, The Netherlands, 2005.
4. Dicky Suryadi and Piotr J. Gmytrasiewicz. Learning Models of Other Agents Using Influence Diagrams. In *Proceedings of the Seventh International Conference on User Modeling*, 223–232, Banff, Canada, 1999.
5. Julita Vassileva, Gordon McCalla, and Jim Greer. Multi-Agent Multi-User Modeling in I-Help. *User Modeling and User-Adapted Interaction*, 13(1-2):179–210, 2003.
6. Ingrid Zukerman and Christian Guttmann. Modeling Agents that Exhibit Variable Performance in a Collaborative Setting. In *Proceedings of the Tenth International Conference on User Modeling*, Edinburgh, Scotland, 2005.