Automated Aiding Strategies for Decentralized Planning with Interdependent Policies

(Extended Abstract)

Martin J. Kollingbaum
Robotics Institute
Carnegie Mellon University
mkolling@cs.cmu.edu

Joseph Giampapa
Robotics Institute
Carnegie Mellon University
garof@cs.cmu.edu

Katia Sycara
Robotics Institute
Carnegie Mellon University
katia@cs.cmu.edu

Timoty J. Norman
University of Aberdeen
t.norman@abdn.ac.uk

Chris Burnett
University of Aberdeen
cburnett@abdn.ac.uk

Daniele Masato
University of Aberdeen
d.masato@abdn.ac.uk

ABSTRACT

We discuss how agents can support collaborative planning activities within coalitions, where coalition partners plan and act according to given policies. In a set of experiments with human test subjects, we investigate how agents can aid such a collaborative planning effort and how effective they are in the chosen support modes.

1. INTRODUCTION

Coalitions are formed to collaborate in the performance of specific tasks. A plan must be constructed to coordinate the actions of the coalition partner for a successful completion of such tasks. This poses a particular challenge, especially in the face of individual goals, self interest and with coalition members having only limited co-training for recognizing and resolving their differences. In a set of experiments, we investigated how agents may influence and coach such a distributed planning effort. In particular, we are interested in a situation where this planning process is influenced by policies – each coalition partner may have to uphold certain normative standards describing their obligations, permissions and prohibition during their collaborative planning activity. With such policies in place, human planners are under pressure to produce high-quality plans while, at the same time, adhering to all their obligations and prohibitions as described by their policies.

2. EXPERIMENTS

We chose a rescue mission as an example scenario and performed a set of experiments to investigate the effectiveness of agents supporting a collaborative planning effort. In this scenario, we assume that there are two parties that form a coalition, a humanitarian relief organization (regarded as "Party A") with the individual goal of rescuing injured civilians from a potentially hostile region, and a military organization (regarded as "Party B") that has to coordinate its military objectives with the evacuation activities. The goal of this coalition is to find a joint plan for transporting as many injured people from a dangerous region to a hospital in the shortest possible time. We investigated the following supporting stragegies or "aiding conditions" for an agent. In the *critic* condition, the agent detects policy violations that may occur during the human planners' communication and planning behavior. If the human is about to commit a communication violation, the agent intercepts the violating message and presents text explaining the policy that the message risks violating. If the human is about to commit a planning violation, the agent behaves in the same way before the plan step is committed to the user's plan. In both situations, the user may decide to either accept the agent's advice, and reformulate their communication or plan step, or to reject the agent's advice, override the intervention, and commit the violating communication or plan step, anyway. In the *censor* condition, the agent still monitors the activities of the human planner, but silently interferes with the communication by deleting offending parts of the exchanged message (or block it completely) in order to avert policy violations. The *critic* agent does not force the user to a particular action, it merely provides advice and suggestions that the user can accept or reject. The *censor* agent intercepts any message that contains policy violations and suppresses it. Only the recipient is informed that the sender's message contained a policy violation and was suppressed. Both conditions are compared to a third control condition – the *unaided* condition, where the test subjects did not have any agent support.

The human planner interacts with the agent via a software infrastructure that supports the communication within the coalition and assists in the planning task. The purpose of this application is to assist the user in their communications and planning with coalition partners, while at the same time being respectful of their organization's policies. It is our hypothesis that the benefit of the agent assistance is noticed through improved performance metrics such as: the number of violations that a user commits in the course of their collaborations, the number of wounded that are rescued, and the overall costs of the individual and combined plans.

3. RESULTS

Cite as: Automated Aiding Strategies for Decentralized Planning with Interdependent Policies, (Extended Abstract), Martin J. Kollingbaum, Joseph Giampapa, Katia Sycara, Timoty J. Norman, Chris Burnett, Daniele Masato, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009),* Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1313–1314

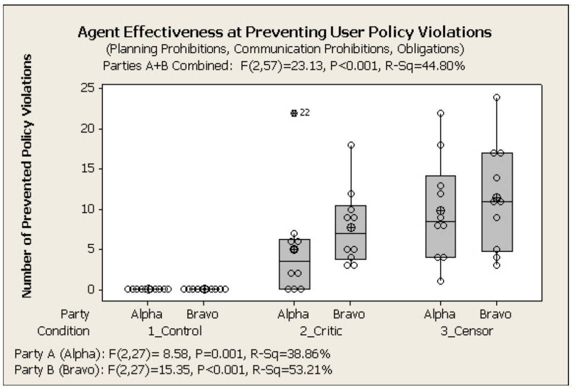

Figure 1: The number of times that a user did not commit a policy violation due to agent intervention.

Figure 1 illustrates the effectiveness of agent interventions at preventing policy violations in the three conditions: "control", "critic" and "censor". It shows the number of times that an agent detected a violation and the user either accepted the agent's feedback or at least did not try to defeat its intervention. There was a minority of human subjects, however, that adjusted the ways in which they used the agents as a result of the type of agent intervention. For example, in the censor condition, some users would try to exhaustively generate-and-test communications for granting clearances or committing to escorts against the censorship of the agent.

Tests of statistical significance were applied to different groupings of the results: (1) Party A across the three conditions, (2) Party B across the three conditions, (3) the combination of Parties A and B across the three conditions and (4) differences among the number of effective interventions per party (no statistical significance). Post hoc tests in SPSS on ANOVA group (3), above, revealed consistent differences ($p < 0.001$) for all three pairwise comparisions: control vs. critic, control vs. censor, and critic vs. censor.

A set of characteristics could be observed about both conditions. 60% of the subjects in both Parties A and B felt that it was necessary to override the critic agent in order to complete their plans. Only one subject out of thirty actually reached the mission objective of treating 100% of the wounded on the first day of the mission without violating any policies, but with the assistance of 13 interventions from the critic agent. Of the two agents, critic and censor, the censor agent was the most effective at preventing policy violations. The Party A subjects in the censor condition were most distracted from their mission objective of treating as many wounded as possible on day 1. We hypothesize that the lack of direct feedback to the user committing the violation may be the cause. Similarly for both parties, the plan costs were greatest in the censor condition.

4. CONCLUSION

The results of these experiments have shown that agents can have a positive impact on the enforcement of policies. We saw that in the unaided experiment condition (no agent monitoring and feedback), individuals would make on average from $7-10$ policy violations, with all individuals making at least three policy violations, per session. We saw that the rate of individuals deliberately violating policies dropped to a median average of one violation per session, with many individuals not making any policy violations at all.

Results from these experiments provide us with feedback that informs the design of future policy-enforcing agents. For example, the censor agent was clearly the undisputed best enforcer of policy. Its enforcement reliability was countered by a significant increase of plan costs and distraction from the mission objectives for Party A. The critic agent immediately flagged policy violations, so it was possible for at least one subject in all three conditions to achieve the perfect plan (e.g. treating all the wounded on day 1), as well as enabling others to have zero mission impacting policy violations. A possible next agent to test would be a critiquing censor agent that: is capable of critiquing plan steps, provides direct feedback to the user that it is censoring, and does so without allowing the censored subject to override it.

5. RELATED WORK

Our concept of a policy is alligned with research into normative systems [2], in particular work on norm-governed agency [5] and Electronic Institutions [4, 6]. In this paper, we expand on work presented in [1] and describe a specific application of normative agents, where agents do not form virtual organisations, but observe the behaviour of human planners within a coalition. By using a rule-based language for encoding policies and, consequently, a rule engine such as JESS (Java Expert System Shell) [3] for processing these specifications, we followed a very similar implementation path as described in [4], in particular implementing the reasoning about policies.

6. ACKNOWLEDGMENTS

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

7. REFERENCES

[1] C. Burnett, D. Masato, M. McCallum, T. J. Norman, J. Giampapa, M. J. Kollingbaum, and K. Sycara. Agent Support for Mission Planning Under Policy Constraints. In *Proceedings of the Second Annual Conference of the International Technology Alliance*, pages 100–107, 2008.

[2] F. Dignum. Autonomous Agents with Norms. *Artificial Intelligence and Law*, 7:69–79, 1999.

[3] E. Friedman-Hill. *Jess in Action*. Manning, 2003.

[4] A. García-Camino, P. Noriega, and J. A. Rodríguez-Aguilar. Implementing Norms in Electronic Institutions. In *Procs. 4th AAMAS*, 2005.

[5] M. Kollingbaum. *Norm-governed Practical Reasoning Agents*. PhD thesis, University of Aberdeen, 2005.

[6] W. W. Vasconcelos. Norm Verification and Analysis of Electronic Institutions. In *DALT 2004*, volume 3476 of *LNAI*. Springer-Verlag, 2004.