Agent Support for Policy-driven Collaborative Planning in Ad-hoc Teams

Martin J. Kollingbaum
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: mkolling@cs.cmu.edu

Joseph A. Giampapa
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: garof@cs.cmu.edu

Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: katia@cs.cmu.edu

Timothy J. Norman
Dept. of Computing Science
University of Aberdeen
Aberdeen, AB24 3UE, UK
Email: t.j.norman@abdn.ac.uk

***Abstract*—The creation of joint plans within teams is a complex task, especially if these teams are formed in an ad-hoc fashion with limited co-training. Team members may have to plan their actions in accordance with a set of regulations or mission policies and by observing planning constraints. For ad-hoc teams operating under time-stressed conditions, this is a difficult task. In this paper, we describe how to construct agents that can support teams in their collaborative planning effort. We show how agents can be integrated into the planning and communication activities of human planners. Agents monitor human planners, reason about their actions and advise them on possible violations of mission policies and planning constraints.**

I. Introduction

In the formation of coalitions for the purpose of performing rescue missions, the planning process for such missions may be complicated by the fact that coalition members have to operate under a set of rules or mission policies that describe what their obligations, permissions and prohibitions are in terms of the actions that can (or have to be) planned and the communication necessary to coordinate planning activities with coalition partners. Taking such a body of regulations into account during planning can amount to a substantial cognitive burden for human planners, especially in producing high-quality plans in time-stressed situations and with members of such coalitions having only limited co-training as a preparation for collaborating in joint missions and forming joint plans.

We assume that mission policies, planning constraints and constructed plans are private information of a coalition member, which may or may not be disclosed to other coalition partners. With that, coalition members find themselves in a situation where they have to form coalitions without agreeing to (and knowing) a common set of policies or "social norms" or each other's private plans and still have to find a way to operate in a collaborative fashion. As coalition partners may depend on each other in performing joint missions or have to clarify they interfere with each others' actions, they have to communicate certain information. This information exchange may pose a particular security risk as information may be disclosed either intentionally or unintentionally and, therefore, may be subject to tight regulations / policies. For a human planner to produce a plan on short notice with these policies in mind, this may pose a huge cognitive burden. In this paper, we describe, how to construct an agent that can assist human planners in a monitoring and advisory capacity. For this, we have to meet a set of challenges: (a) mission policies have to be represented in a form that can be processed by a software agent, (b) such an agent must be able to keep track of the "normative state" of the human planner – what are the planner's current obligations, permissions and prohbitions, and (c) the agent must be able to reason about the current actions of the human planner, so that it can provide feedback about violations of policies, and advise the human planner about alternative courses of action.

II. Scenario

We assume that two parties have to form a coalition for the purpose of performing rescue missions. A humanitarian relief organization with the individual goal of rescuing injured civilians from a potentially hostile region cooperates with a military organization in order to achieve this objective. Both coalition partners have to coordinate their activities, deploying particular resources to rescue wounded or pursue military objectives.

III. Agent Support

We introduce personal agents to support and advise human planners in their planning and communication activities. These personal agents are directly involved in a human planner's activities by monitoring the communication acts of a human planner as well as planning activities. The set of policies or *norms* relevant to human planners determine their "normative" state or *normative position*. Agents, as we utilize them, have to maintain a representation of this normative state for the monitored human planner. Currently, we investigate two main aiding strategies of an agent:

- in the *critique* condition, the agent detects policy violations that are incurred by the human planners in their communication and planning behavior. In case of a violation, the agent either (a) intercepts the sending of a message or (b) interrupts the planning of actions by informing the human planner about the set of policies that become relevant due to these intended and violating actions – the planner can then decide whether to adhere to such an advice or to ignore the agent and intentionally violate a policy;

- 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 messages (or blocking them completely) in order to avert policy violations; in that case, the *receiver* is informed that a message is either truncated or completely censored.

The difference between the two types of agents is in their policy-related feedback to the human planner and their subsequent interaction. The critic agent, besides reasoning about policies, also monitors plan steps committed by a human planner and reasons about the effect of policies on planned actions. The censor agent, on the other hand, is not concerned with effects of policies on planned actions, it only interecpts and forbids the transmission of messages that contain policy violations. The *critique* agent does not force the user to a particular action, it merely provides advice and suggestions, which the user can accept or reject.

IV. Policies

Policies are given to the human planner in a verbalized form. At the same time, they have to be represented in a form that can be processed by an agent. We specify policies in the following way. If we define the set $Expr$ as the set of all possible well-formed formulae comprising first-order predicates over terms (constants, variables and the operators $\wedge$, $\vee$ and $\neg$, then a policy can be defined in the following way:

Def. 1: A policy, expressing an obligation, permission, prohibition is a tuple $\langle \nu, \rho, \varphi, a, e \rangle$, where

- $\nu \in \{\mathcal{O}, \mathcal{P}, \mathcal{F}\}$ is a label indicating whether this is an obligation, permission or prohibition
- ρ is a role identifier for a norm addressee
- φ describes the action regulated by this policy
- $a \in Expr$ is the activation condition
- $e \in Expr$ is the expiration condition

This definition displays in a simple fashion the elements that characterize an implementation of our policies – they are ascribed to a specific role and are activated/de-activated under certain conditions. The policies themselves exist in two forms, (a) formulated in simple "IF ... THEN ..."-statements that are given to human planners, and (b) as a set of rules (implemented in the Jess expert system shell language [1]), expressing their activation/de-activation, in order to allow agents a processing of these policies and the reasoning about their current activation state.

V. Planning Constraints

Besides policies, that express "social" regulations or norms of behaviour for coalition partners in a mission, we also have to consider basic operational or planning constraints. These constraints determine the *feasibility* of a plan operation. In general, we consider several types of planning constraints. *Physical constraints* capture the physical conditions under which coalition partners may employ particular actions, such as passability of roads, transport capacities of vehicles, etc. With *routing / planning / scheduling constraints*, we capture, for example, that a vehicle cannot be dispatched to more than one city in the same day. *Resource constraints* capture situations of resources being deployed to multiple destinations at the same time. With *quality-of-plan constraints*, we can capture quality critera such as, for example, "treat all wounded on the first day of a two-day mission". Explicit *policy constraints* test policy violations related to planning activities.

VI. Related Work

Policy-based reasoning (e.g. Ponder2 [2]) is used to regulate access to shared resources or to enforce certain properties (quality of access / service / data etc.). Our concept of a policy is inspired by previous work [3] and is aligned with research into normative systems, in particular work on norm-governed agency [4] and Electronic Institutions [5], with [6] discussing how human norms can be represented in a form that allows agents within e-institutions to process them.

VII. Conclusion

In this paper, we discussed difficulties in establishing joint plans within coalitions in the face of self interest, individual goals and diverse policies of the coalition members. As we showed, detailed and, sometimes, even conflicting policies have to be dealt with in practice, when coalitions try to engage in collaborative planning. We therefore advocate agent support for policy-based planning activities within coalitions. In this paper, we demonstrated how agents can be integrated into the dialogical process of human planners establishing a collaborative plan.

Acknowledgement

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.

References

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

[2] M. Charalambides, P. Flegkas, G. Pavlou, A. Bandara, E. Lupu, A. Russo, N. Dulay, M. Sloman, and J. Rubio-Loyola, "Policy Conflict Analysis for Quality of Service Management," in *6th IEEE Workshop on Policies for Distributed Systems and Networks (Policy 2005)*, 2005.

[3] M. Kollingbaum, "Norm-governed Practical Reasoning Agents," Ph.D. dissertation, University of Aberdeen, 2005.

[4] F. Dignum, "Autonomous Agents with Norms," *Artificial Intelligence and Law*, vol. 7, pp. 69–79, 1999.

[5] W. W. Vasconcelos, "Norm Verification and Analysis of Electronic Institutions," in *DALT 2004*, ser. LNAI. Springer-Verlag, 2004, vol. 3476.

[6] J. Vazquez-Salceda, H. Aldewereld, D. Grossi, and F. Dignum, "From Human Regulations to Regulated Software Agents' Behaviour: connecting the abstract declarative norms with the concrete operational implementation," *Artificial Intelligence and Law*, vol. 16, no. 1, March 2008.