Overview of HANA: a Human-Aware Negotiation Architecture*

Angela Fabregues and Carles Sierra

Artificial Intelligence Research Institute (IIIA-CSIC)
Campus Universitat Autònoma de Barcelona, 08193 Bellaterra, Catalonia, Spain
{fabregues,sierra}@iiia.csic.es

1 Introduction

This position paper is an overview of the journal article [1] where we propose HANA, a software architecture for agents to sign agreements about joint plans of action in a multiagent system (MAS) with repeated interactions. Negotiations are bilateral and may involve humans. We assume that there are specific time deadlines for action execution and that the agents have to negotiate while exploring what set of actions to perform.

Realistic problems involving humans use to be of the incomplete information type and have a huge space of solutions. That means that the space is large enough and the negotiation time short enough to preclude looking for the optimal solution. That means that any architecture for this type of negotiation problem needs to give the means to look for *g*ood-enough solutions.

2 Architecture

HANA's architecture is graphically represented in Figure 1. It provides an interface that situates HANA agents in their environment via some sensors and actuators. The sensed information is stored in the world model as beliefs. The agents have a BDI model extended with emotions that are used to update the agent's desires and to influence the generation of intentions from beliefs and desires. These motives are all graded [2] in order to cope with the unavoidable world's uncertainty. In this architecture action planning is not part of the BDI model.

Evaluating plans and options (partial plans that can be negotiated) requires a good model of the world. To be effective, this model has to include the actions to be performed by the other agents. This is so, as the outcome of the actions of an agent depends on the actions performed by the rest. Reducing the uncertainty on what the others will do is thus crucial for the agent success and this can only be done by reaching commitments of others through negotiation. Every signed agreement adds commitments to agents that constrain their possible actions and thus reduce the world's uncertainty for those that know those commitments.

* AT2012, 15-16 October 2012, Dubrovnik, Croatia.

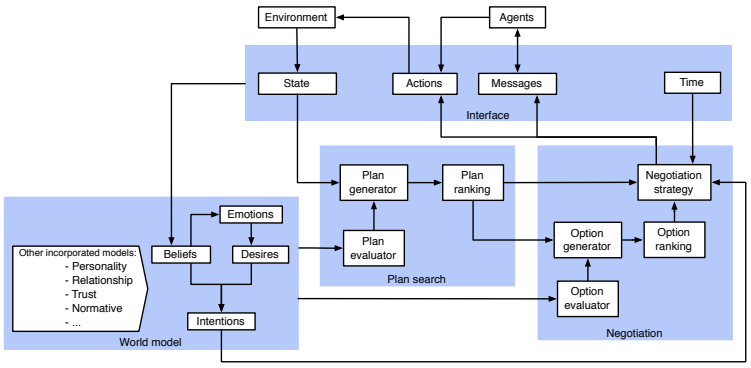

Fig. 1. HANA architecture. Arrows represent data flows. Coloured boxes represent modules and white boxes are components of those modules.

Moreover, in realistic settings the longer it takes for an agent to decide what to propose the less probable it is that the proposal will be accepted. As time goes by, agents reach agreements increasing the amount of commitments and reducing the set of options that are compatible with previous commitments. Therefore, the architecture must allow a swift start of the negotiating dialogues.

HANA agents interact from the very beginning, and make proposals while they keep on searching for better plans. The execution of an agent consists of several concurrent processes. We use a new search&negotiation technique where plan search and negotiation go hand in hand: the former generating proposals and the later pruning the search space towards better plans. The search is done by an anytime algorithm that provides a ranking of plans that are then used to generate a ranking of options to be negotiated. Every signed agreement provides a set of commitments that prune the search space of feasible plans. This new technique is crucial for bilateral negotiations in MAS with a huge solution space and with time constraints.

HANA's architecture is modular and facilitates the extension of basic HANA agents with more sophisticated behavioural models that refine the generation of plans and options.

Acknowledgments

Research supported by the Agreement Technologies CONSOLIDER project under contract CSD2007-0022 and INGENIO 2010, and the COST Action, IC0801.

References

1. Fabregues, A., Sierra, C.: Hana: a human-aware negotiation architecture. Decision Support Systems (In press.)
2. Casali, A., Godo, L., Sierra, C.: A graded bdi agent model to represent and reason about preferences. Artificial Intelligence **175** (2011) 1468–1478