Extending the MOISE model with adjustable autonomy

Artur Vidal Maia, Jaime Simão Sichman

[1]Laboratório de Técnicas Inteligentes (LTI)
Escola Politécnica – Universidade de São Paulo – USP – SP – Brazil

artur.vidal@gmail.com, jaime.sichman@poli.usp.br

Abstract. *Having autonomy and being part of an organization can be opposite concepts. Autonomy lets an agent take his own decisions, while sometimes, being part of an organization constraints these decisions to the ones allowed by the organization. This works presents an idea to combine these aspects, by allowing an agent to have autonomy even if it take part of an organization. We present different types of roles and goals to tackle this problem, characterizing what is called adjustable autonomy.*

1. Introduction

One of the main issues in MAS are social organizations, which consist of the set of relations and interactions between cooperative agents to achieve an objective. These relations and interactions are usually modeled after real world organizations. In order to represent these social organizations in a computer system, various models have been proposed.

The study of organizations is mostly present in Multi-Agent Systems as a way to optimize the behaviour of a given system, and different organizations can provide different results. The work described in [Franco and Sichman 2014] has proposed different organization specifications to solve a same problem, namely the *Agent on Mars scenario* [Behrens et al. 2012]. These different organizations were tested in a common set of environments, and their performances were compared with respect to some environment patterns. On the other hand, it is well-known in management sciences that the problem of identifying the *best* organization for a certain scenario is a very hard problem. Normally, real organizations evolve in time, leading to what is called reorganization process. As shown by [Hübner et al. 2004], reorganization is a very difficulty process and even if we can detect which parts of the organization are failing, we need to consider many aspects for proposing useful corrections.

For a concrete case, let us consider, as a dynamic environment, an office building where there is a population of humans working with robots. These robots execute general services such as garbage collection, mail distribution, etc. However, in emergency situations, such as an earthquake, the robots assume fire combat functions. The robots are members of a multi-agent organization and its valid to suppose that the pre-defined organization for the daily environment is not the best one to handle the emergency scenario. A global reorganization approach for this problem should detect this emergency situation and change the current organization for the new scenario.

We believe that an organization which supports *adjustable autonomy* is a better approach than a global reorganization approach. Adjustable autonomy allows a smooth transition from one environment to the other since the agents may detect the changes by themselves. Note that in the global reorganizing approach, we still have the problem of designing an organization for each case, which can not be known beforehand in the case of dynamic environments.

That is where the concept of adjustable autonomy can be used to better solve this problem. Adjustable autonomy can be characterized as the possibility of change in the decision making process of an agent by having it either imposed by an external entity (the organization) or defined by the agent itself. Previous work [Barber et al. 2000] has showed that having adjustable autonomy in a MAS can greatly improve the performance of the system by allowing the decision making process to change priority between local and global knowledge. This is more important in dynamic environments where a centralized coordination mechanism might fail as information becomes outdated or imprecise. Using adjustable autonomy, we face the problem of designing an organization from another point of view. Instead of changing its specification, we look at how to design an organization whose components can adapt to the environment where they are situated; more precisely, we intend to propose an organization model which *explicitly* allows its agents to adapt.

The main motivation of this work is that we expect that an organization with adjustable autonomy improves its performance by allowing better agents' adaptability while still maintaining some control over them. Autonomous agents can create and change their behaviors without changing the organizational specification, enabling them to react and adapt more easily, without violating the organizational constraints. However, autonomy is not a simple concept and currently there are various definition for different types of autonomy. We have decided to use in our work the notion of planning autonomy, as detailed in the sequence.

2. Adjustable Autonomy

Autonomy is the ability that an agent has to take his own decisions unconstrained by external elements. In [Falcone and Castelfranchi 2001, Falcone and Castelfranchi 1999], the authors propose the concept of *adjustable autonomy*, enabling to define different degrees of autonomy. Adjustable autonomy is characterized by two important attributes: Control and Independence [Barber et al. 2000].

Control is the attribute which characterizes the importance level of the agent in the decision making mechanism. In multi-agent systems, it is not uncommon to have a collective high-level decision making mechanism which takes into account the participating agents' preferences. An agent with high level of control means that he takes his own decisions, while an agent with a low level of control means that the decisions are mostly taken externally.

Independence is the attribute which shows the number of options an agent has to achieve its goals. High independence means that an agent has many possibilities to achieve a certain goal, while low independence means that an agent has one or just a few possibilities to achieve a goal. Figure 1 shows how control and independence are related to an agent. In the left agent, agent 3 has a low control level since agents 1 and 2 interfere with his decisions. On the other hand, he has a degree of independence as he can achieve his goal by using different plans.

Adjustable autonomy has been successfully used as a coordination mechanism in multi-agents systems evolving in dynamic environments [Van Der Vecht et al. 2008].

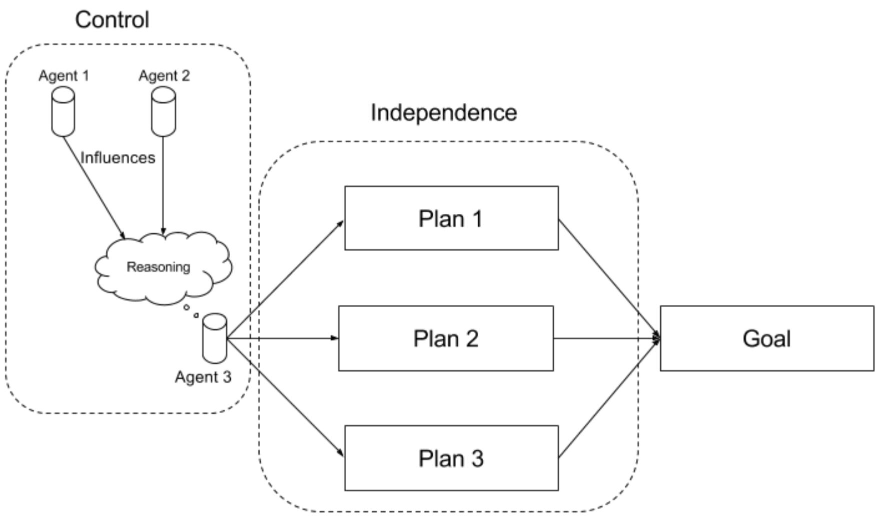

Figure 1. An agent's adjustable autonomy dimensions

The possibility of changing independence and control during runtime can make a system perform better by allowing agents to quickly adapt to different situations by having higher independence and control; on the other hand, agents may save time by having their autonomy diminished when there is already an agent with a plan to solve the problem. An agent's autonomy adjustments can be triggered from different sources, like another agent, the organization he takes part in, or the agent itself.

Our proposal to model adjustable autonomy use the concept of role types and goal types [Van Riemsdijk et al. 2008] to implement planning autonomy.

3. Proposal and Methodology

We are using the MOISE organizational model [Hübner 2003] as a buillding block for our work. It is composed by three specifications: structural (SS), functional (FS) and deontic (DS) specifications. Concerning goals, the organization decomposes high level goals in sub-goals (FS) that are associated to the agents playing certain roles in groups (SS), and this association is defined in terms of obligations or permissions (DS). As it can be seen, the main objective of an organization is to represent in an indirect way how the agents will cooperate. An organizational model is set of norms whose instantiation results in a concrete and specific organization entity that should solve a problem. The paradigm is that the agents assume roles in the organization, and try to accomplish the associated goals through missions assigned to these roles.

The functional specification (FS) is used to determine processes and workflows. While this specification can be used to create relations between goals, it does not contain more information about how goals can be achieved or what are these goals in terms of logical world states. So, effectively these goals can be treated as sort of checkpoints for the agents to tell the organization that they reached a certain state, rather than a concrete description about what needs to be reached. This kind of description is usually sufficient, since the organization is considered as a normative system and hence agents must adopt specific goals and act until they "achieve" them; however, this leads to situations where agents may tell the organization that they achieved a goal while they actually have not, or even their definition or interpretation of an organizational goal might not be the one made up by the knowledge engineer who designed the organization.

Some previous work [De Brito et al. 2015] proposed a way to link the normative description of an organization goal with a concrete one, which helps to unify what a goal means to all agents. We will use these concrete descriptions of goals as a basic building block to implement agents with planning autonomy. In classical planning, goals are defined as desired states of the world which a planning algorithm must achieve through the successive use of actions from a starting state. Using a tree representation, states of the world are nodes and the actions are edges. A particular case of planning which considers task decomposition is HTN [Nau et al. 2003]. In HTN, these planning trees have a finite set of specific possible patterns for the branches. These patterns are called methods and they represent intermediate state constraints which should be reached. Since both HTN planning and the functional specification of the MOISE have a goal decomposition tree, we intend to use the HTN formalism to describe parts of the functional specification. Some previous work addressed the integration of automated planning in a multi-agent framework [Cardoso 2014], but the authors proposed to use generated plans from a multi-agent planning algorithm as an initial organization specification of a MAS. We will incorporate such idea as a way to improve the current MOISE model by allowing it to interact with planning agents.

In the current MOISE model, agents have complete freedom to reason about their goals: they can plan or choose between any predefined plans to satisfy them. However, even if in certain cases planning autonomy is positive (for instance, when an expert agent can submit plans to the organization so that other agents can use his expertise), in some other cases the imposition, by the organization, of a certain predefined plan saves time for the global system (since the agent does not need to reason about the current goal). Currently, the functional specification (FS) of the model does not allow theis difference to be defined: a leaf in a goal tree may mean either an atomic action or a subgoal whose achievement must be autonomously defined by the role playing agent.

Our proposal is to improve the current MOISE model by including a definition of two different types of goals. A first one, called *Plan-Goal*, consists in creating a relation between the current MOISE goal and a resulting plan from a planning algorithm. This is intended as a way to allow current MOISE goals to be more tightly related to a specific plan, thus constraining more precisely the agent's autonomy. On the other hand, the second type, called *State-Goal*, represents a world state that the organization expects to be achieved. This second type of goal allows agents to reason about them and to try to achieve them as they like: in other words, agents may autonomously decide how to achieve them. Hence, agent's autonomy are expressed through the relations of the agents's roles with the different types of goals: the more State-Goals he has, the more autonomous he his. Finally, adjustable autonomy may be carried out by an organization when sending either a Plan-Goal or a State-Goal to the same agent, in different situations, based on the agent's current mission.

In parallel, we need to refine MOISE's Structural Specification by defining two types of roles: *Operational* and *Management* roles. The fist type is intended to be destinated to those agents that should not have autonomy, i.e. the organiztion defines how they will achieve their goals. On the other hand, agents playing management roles are those intended to have adjustable autonomy: they may receive both Plan-Goals and State-Goals to be achieved.

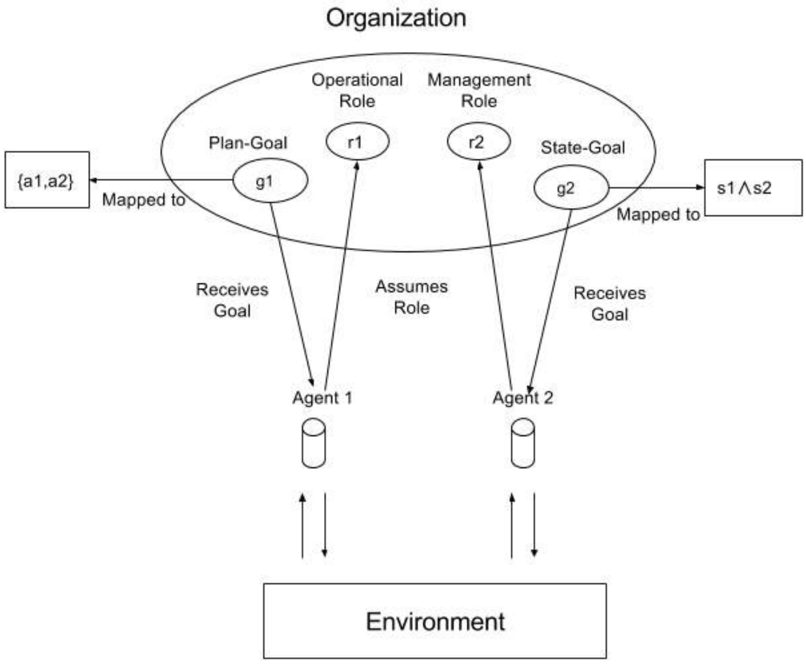

Figure 2. An organization with different role and goal types.

Figure 2 illustrates our proposal, where both types of roles and goals are part of an example of organization specification. On the left side, Agent 1 assumes an operational role $r1$, to which then organization defines a Plan-Goal $g1$ to be followed, and this plan is mapped to a sequence of actions $\{a1, a2\}$. On the right side, we notice Agent 2 assuming a management role, which is associated to a State-Goal $g2$: in this case, the agent is allowed to create his own plan of action to achieve the desired world state $s1 \wedge s2$. In this example, both agents have different degrees of autonomy, and consequently different degrees of control and independence.

4. Conclusions and Further Work

We propose in this work to incorporate the concept of adjustable autonomy in the MOISE organizational model. For this, we define two types of goals (Plan-Goal and State-Goal) and roles (Operational and Management roles) to express different control and independence levels.

Regarding evaluation, we intend to run an experiment to compare the results obtained by two different organizations, one with and another without adjustable autonomy. This experiment should have as a main characteristic a dynamic, partially observable environment. Environments with these characteristics are expected to produce better results when using adjustable autonomy rather than the original MOISE model. There is a number of popular scenarios with these characteristics in the literature, such as the *Agent on Mars scenario* [Behrens et al. 2012].

References

Barber, K. S., Martin, C. E., Reed, N. E., and Kortenkamp, D. (2000). Dimensions of adjustable autonomy. In *Pacific Rim International Conference on Artificial Intelligence*, pages 353–361. Springer.

Behrens, T. M., Dastani, M., Dix, J., Hübner, J., Köster, M., Novák, P., and Schlesinger, F. (2012). The multi-agent programming contest. *AI Magazine*, 33(4):111–113.

Cardoso, R. C. (2014). Integrating automated planning with a multi-agent system development framework. Master's thesis, Pontifícia Universidade Católica do Rio Grande do Sul.

De Brito, M., Hübner, J. F., and Boissier, O. (2015). Coupling regulative and constitutive dimensions in situated artificial institutions. In *European Conference on Multi-Agent Systems*, pages 318–334. Springer.

Falcone, R. and Castelfranchi, C. (1999). Levels of delegation and levels of adoption as the basis for adjustable autonomy. In *Congress of the Italian Association for Artificial Intelligence*, pages 273–284. Springer.

Falcone, R. and Castelfranchi, C. (2001). The human in the loop of a delegated agent: The theory of adjustable social autonomy. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 31(5):406–418.

Franco, M. R. and Sichman, J. S. (2014). Comparing and evaluating organizational models: a multi-agent programming contest case study. In *International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems*, pages 182–196. Springer.

Hübner, J. F. (2003). *Um modelo de reorganização de sistemas multiagentes.* PhD thesis, Universidade de São Paulo.

Hübner, J. F., Sichman, J. S., and Boissier, O. (2004). Using the\ mathcal {M} oise+ for a cooperative framework of mas reorganisation. In *Advances in artificial intelligence–SBIA 2004*, pages 506–515. Springer.

Nau, D. S., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., and Yaman, F. (2003). Shop2: An htn planning system. *J. Artif. Intell. Res.(JAIR)*, 20:379–404.

Van Der Vecht, B., Dignum, F., Meyer, J.-J. C., and Neef, M. (2008). A dynamic coordination mechanism using adjustable autonomy. In *Coordination, Organizations, Institutions, and Norms in Agent Systems III*, pages 83–96. Springer.

Van Riemsdijk, M. B., Dastani, M., and Winikoff, M. (2008). Goals in agent systems: a unifying framework. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 713–720. International Foundation for Autonomous Agents and Multiagent Systems.