Adaptive Autonomy for Agile Task Coordination

Martijn Neef[1] and Bob van der Vecht[1,2]
[1] Command & Control and Information Management Department,
TNO Defence, Security and Safety, The Hague, The Netherlands.
[2] Department of Information and Computing Sciences, Utrecht University, Utrecht.

martijn.neef@tno.nl / bob.vandervecht@tno.nl

1. INTRODUCTION

Achieving high levels of agility and resilience in upcoming military organizations will require new ways of thinking about command and control. As traditional military organizations are gradually being equipped for new types of network-centric missions, it becomes obvious that we need rethink coordination strategies. In future arenas, there will be many more parties involved and a much less transparent chain of command. We will need to rely more on distributed processes and accept that traditional centralized command and control strategies will not lead us to agile capabilities.

Our present work on adaptive autonomy in agent systems might provide some interesting perspectives on agility and resilience in NEC environments. We have been researching the topic of autonomy in multi-agent systems, and in particular the relationship between autonomy and coordination. In any multi-actor environment, there is an inevitable trade-off between achieving global coordination of activities and respecting the autonomy of the actors involved. This is also clearly an issue in NEC-oriented organizations. If decision making processes and operational tasks are distributed over many parties, respect of autonomy becomes an increasingly important issue. We are working on decision making models that respect the agent's own autonomy, but at the same time take organizational roles and operational conditions into account.

2. ADAPTIVE AUTONOMY

Autonomy is an important aspect of artificial agents. It is usually regarded as one of the defining features of an agent [Jennings2000, Castelfranchi1995]. Agents have control over both their internal state and behaviour. Agents act goal-oriented, and exert their autonomy to reach their objectives. In multi-agent systems, there is communication and coordination of activities between agents. Coordination implies that agents can influence each other, and possibly make demands that may affect each other's degree autonomy. Since, by definition, agents need to be in control of their own internal state and behaviour, the question rises how agent decision making is actually impacted by external influences. How can an agent maintain control over his own autonomy, but at the same time cooperate with other agents to achieve coordination?

The traditional way to achieve coordination is by developing a top-down coordination mechanism. The designer of a multi-agent system specifies the tasks and interaction mechanisms that the agents will follow. The rules of the coordination mechanism are embedded in the decision-making process of the agent. This allows the agents to jointly find a correct division of labour. One can argue that such agents are not truly autonomous, since they can only behave in line with commands that are set forth at design-time. The agent has no means to enforce its autonomy, and cannot pro-actively deviate from plans. Of course, in many systems, this is a desirable feature: we want the system to do what it was designed for. In some cases, however, it might be favourable to grant agents a degree of sovereignty. One could think of situations where standard procedures fail, and agents are left to their own devices to create alternative plans. For instance, if the chain of command collapses because of communication breakdown, deployed agents need to be able to take matters into their own hands. In other words, they need to adapt their autonomy from being commanded to self-ruling.

In our perspective, autonomy is about the level of independence of decision making. The degree of autonomy of decision making can be defined as the degree of intervention by other agents on the decision making process of one agent [Barber2001]. An agent that is heavily influenced by other agents in its decision making is displaying obedient behaviour. An agent that does not allow any external influence in its decision making is being ultimately independent. By altering the amount of external influence on its decision making, an agent can adapt its autonomy, and show adjustable autonomy. An agent that switches between different levels of autonomy of its decision-making shows adjustable autonomy. In this fashion, agents can actively select the level of autonomy that best fits the circumstances.

3. CONTROLLING EXTERNAL INFLUENCES

We have developed a reasoning model that gives an agent control over external influences, and, therewith, guarantees the autonomy of the agent [Vecht2007a]. In the reasoning process we distinguish a phase for event-processing and a phase for decision-making, as shown in figure 1. The event-processing phase gives the agent control over its autonomy. The decision phase focuses on the decision on action.

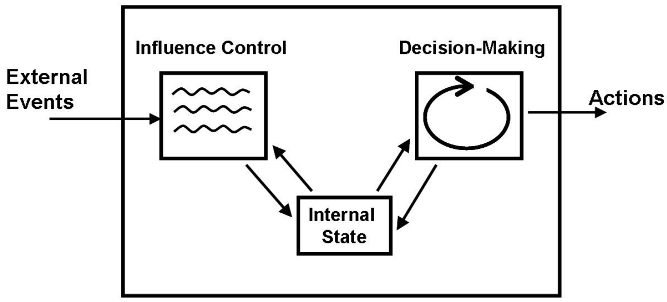

Figure 1: Schematic reasoning model of an adaptive autonomous agent

The model uses reasoning rules to decide on adoption or rejection of certain influences. These reasoning rules weigh knowledge from the agent's internal state against locally available information, and thus may permit or bar external events from influencing the agent's beliefs.

4. HEURISTICS FOR INFLUENCES CONTROL

What knowledge should be taken into account by the reasoning rules? We have explored several heuristics that seem appropriate to control external influences. One of the heuristics is *relevance of information*. If an agent can determine the relevance of information with respect to a certain goal, it can focus itself on a specific type of information, or prevent itself from information overload by filtering incoming information on relevance. Information relevance is important for influence control. Another related heuristic is the *state of mind* of the actor. An actor will react differently when it is busy than when relaxed, or when it feels endangered. We cluster such heuristics as *self knowledge*.

Self knowledge creates heuristics for an agent to determine how it is influenced. The agent should also be able to control *by whom* it is influenced. The reasoning rules for event control can use knowledge about the existing organization or about the agent's social context. Can the sender of a message be trusted? Does a request originate from a superior or from an unfamiliar source? Agents can achieve coordination by allowing influence on the internal state based on social and organizational knowledge. An organizational model describes the roles and relations between actors, and specifies behavioural rules in terms of norms. The modular approach in the agents reasoning model provides a mechanism to adopt organizational rules into the decision-making process. This can be done dynamically by changing contracts at runtime [Vecht2008].

One can think of several other reasons to allow or disallow influence on the internal state in a certain context. A specific coordination type puts requirements on the environment, such as availability of communication or information resources. Critical changes in the environment can be used to determine the proper level of autonomy, and thus provides another important heuristic for influence control: *environmental knowledge*.

Table 1 summarizes the above types of knowledge that can be used in reasoning rules for influence control. This is obviously not an exhaustive list, but these three main types of knowledge seem to capture more relevant factors.

Table 1. Examples of meta- knowledge for influence control

Type of knowledge	Examples
Self knowledge	Relevance of information State of mind (e.g. busy, danger)
Organizational/Social knowledge	Relation to information source Can the source be trusted?
Environmental knowledge	Availability of communication Availability of information sources

5. EXPERIMENTS

For our research, we make use of artificial agents. We created a simulation environment in which artificial actors perform coordinated tasks. In this virtual world, we deploy a group of firefighters that collaborate to extinguish fires that appear at random locations. There are two specific roles in this organization: *coordinator* and *firefighter*. The coordinator makes a global plan and tells the firefighters which fire they should extinguish. Therefore, the coordinator has a global view of the whole world. The firefighters perform the actual tasks in the world; they move to a fire location and extinguish the fires, but they, of course, only have local views. There is a hierarchical relation between the two roles, the coordinator is superior of the firefighters and can send orders to the firefighters, which fire they have to extinguish.

We want to show dynamic coordination within this organization. We achieve this by changing the autonomy level of the decision-making process of the firefighters. We have constructed firefighters that show adjustable autonomy. They are at certain moments disobedient to the commands of the coordinator and at other moments, they follow the orders, depending on their local beliefs. We have implemented the following rules for event processing:

- *IF* command from coordinator *THEN* follow command
- *IF* I am in danger *THEN* ignore commands and follow own goals
- *IF* no communication *THEN* follow own goals

The rules ensure that the agents follow the commands, but if there are no commands, they will pursue their goal using local observations. Also, in case of danger the agent will take care of its own safety. The organization switches between different coordination mechanisms. The responsibility of the choice for coordination type is in hands of the actors themselves. We performed some experiments to evaluate the model in practice and assess the effects of dynamic coordination on overall performance. We were able to show that the reasoning model succeeds in adapting individual agent autonomy under changing circumstances, but at the same time leads to coordinated actions among the agents. For detailed results, we refer the reader to [Vecht2007b].

6. APPLICATION IN NEC ENVIRONMENTS

How does the above model translate to operational environments? The most evident use is in supporting coordination activities in multi-party environments. Agile and resilient behaviour demands dynamic coordination capabilities. Task and resource allocation quickly becomes a demanding challenge in joint and combined NEC environments because of individual constraints and demands. Artificial agents could support this process by acting as *proxies*: mediating representatives for all parties involved.

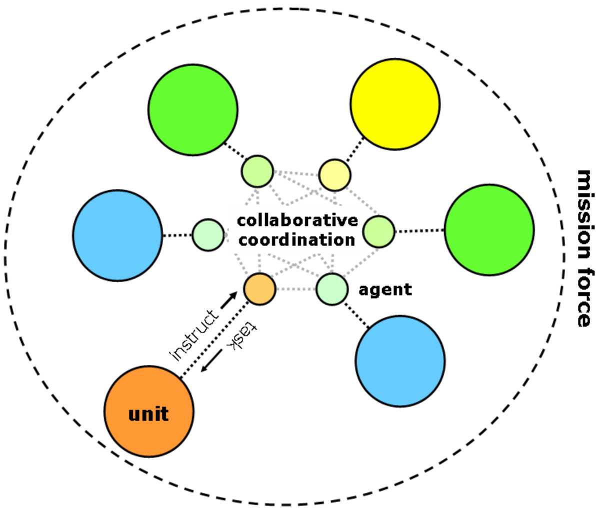

Figure 2: Using mediating agents for collaborative coordination

Such proxy agents could support mission planning and resource sharing, and make it easier to respect individual constraints and policies. Each force member can instruct its proxy agent by tuning the various attitudes to influences. For example, a unit may not be allowed to engage in offensive measures because of local rules of engagement, but may be allow its resources to be used in defensive actions. It could instruct its representative agent accordingly by configuring its openness to external influences. The agent would decide to filter out requests for offensive capabilities, but join the coordination process when dealing with defensive goals. It would use its delegated autonomy to actively accept of refuse requests. In a hectic conflict, there may not be enough time to deal with such individual constraints or resolve potential conflicts through ordinary communication. Artificial agents may help to cope with the dynamics of multi-party collaborations and improve agility. When the organization changes structurally because of leadership changes or the arrival of extra units, the embedded organizational knowledge in the agents

Such autonomy-related configurations could be further facilitated by using meta-reasoning models. In a meta-reasoning model, the agent does not just reason about particular external events, but also over the relationship between various goals and information. This approach gives us a way to use prioritisation of decisions, and in effect construct ‘attitudes’. For example, given a certain operational event, the agent’s reasoning model may conclude that, based on internal knowledge and the agent’s attitude, it prioritises to force protection over self-defence. This means that the agent has received information that it is relevant for the success of two goals (self-defence and force protection), but that it actively chooses to let only the force-protection rule succeed. It blocks the information for the self-defence rule that would lead to retreat.

Because of the modularity of the autonomy model, it is relatively easy to develop and implement such meta-reasoning models and their corresponding attitudes. In practice, such attitudes could serve as a way to implement doctrines.

7. CONCLUSIONS

In this short paper, we briefly introduced our work on adaptive autonomy in agent systems. We have developed a decision making model for artificial agents in which coordination can be defined in terms of organizational norms and rules, but that also guarantees autonomy. We have given the agent capabilities to adapt its openness to external influences, so it can change its own level of autonomy. We distinguish several types of external influences that may impact decision making, namely environmental events, information relevance and organizational rules.

We believe that our approach addresses some fundamental challenges in the progress towards higher NEC maturity levels. Agility and self-synchronization can only be achieved when participants in a NEC organization have adopted practical methods to manage their autonomy. We recognise the essence of having local autonomy, but we also recognise the necessity of coordinated activities. Our model shows that it is possible to relate autonomy and global coordination to each other, and define simple mechanisms that enable adaptive behaviour. The general concept of this model might be used to understand and facilitate task coordination in NEC environments. For instance, networked parties might interact through the use of mediating agents, that represent a party, and guards its autonomy. There will be many practical issues in implementing such a model, but it may inspire NEC developments, and bring about new perspectives on autonomy in collaborative environments.

8. REFERENCES

[Barber2001] - Barber, K.S., Martin, C.E. (2001), Dynamic adaptive autonomy in multi-agent systems: Representation and justification. *International Journal of Pattern Recognition and Artificial Intelligence 15(3)*, pp. 405-433.

[Castelfranchi1995] - Castelfranchi, C. (1995), Guarantees for autonomy in cognitive agent architecture. Intelligent Agents (890), pp. 56-70.

[Jennings2000] - Jennings, N.R. (2000), *On agent-based software engineering*. *Artificial Intelligence* 117(2), pp. 277-296.

[Vecht2007a] - Vecht, B. van der, Meyer, A.P., Neef, R.M., Dignum, F.P.M., Meyer, J-J.Ch. (2007), Influence-Based Autonomy Levels in Agent Decision-Making. In Coordination, Organizations, Institutions, and Norms in Agent Systems II, Lecture Notes in Computer Science, Volume 4386, Springer Verlag, pp. 322-337.

[Vecht2007b] - Vecht, B. van der, Dignum, F.P.M., Meyer, J-J.Ch., Neef, R.M. (2007), A Dynamic Coordination Mechanism Using Adjustable Autonomy. In Proceedings of the MALLOW 2007 Workshop on Coordination, Organisation, Institutions and Norms (COIN@MALLOW 2007), 3-4 September 2007, Durham, UK, pp. 169-180.

[Vecht2008] - Vecht, B. van der, Dignum, F.P.M., Meyer, J-J.Ch., Dignum, V. (2008), *Organizations and Autonomous Agents: Bottom-up Dynamics of the Coordination Mechanism*. In Proceedings of the AAMAS 2008 Workshop on Coordination, Organisation, Institutions and Norms (COIN@AAMAS 2008), , Estoril, Portugal, 12-16 May 2008, Estoril, Portugal (forthcoming).