

RELAXATION OF SOCIAL COMMITMENTS IN MULTI-AGENT DYNAMIC ENVIRONMENT *

Jiří Vokřínek, Antonín Komenda, Michal Pěchouček
Gerstner Laboratory - Agent Technology Center
Department of Cybernetics
Faculty of Electrical Engineering
Czech Technical University
Technická 2, 16627 Praha 6, Czech Republic
{jiri.vokrinek, antonin.komenda, michal.pechoucek}@agents.felk.cvut.cz

Keywords: Social Commitments, Multi-Agent Planning, Relaxation, Uncertainty.

Abstract: The role of social commitments in distributed, multi-agent planning and plan execution will be discussed in this article. We argue agents' capability to reason about the actions in the form of social commitments directly improving robustness of the plans in dynamic, multi-actor environment. We focused on relaxation decommitment strategy, targeted specifically to the time interval in which the agent agrees to accomplish the commitment. We will discuss how changes of this interval affect the plan execution and how the potential changes of this interval can be represented in the commitment itself. The value of the use of social commitments in planning in dynamic, multi-actor environment has been documented on a series of empirical experiments.

1 INTRODUCTION

Cooperation between intelligent agents is usually established by means of negotiation resulting in a set of obligations for the participating agents that lead onwards to achievement of a common goal agreed to by the agents. Wooldridge and Jennings formalize the obligations by describing the cooperative problem solving by means of *social commitments* (Wooldridge and Jennings, 1999) - the agents commit themselves to carry out actions in the social plan leading onwards to achievement of their joint persistent goal (Levesque et al., 1990).

While classical planning algorithms produce a series of partially ordered actions to be performed by individual actors, we propose an extension of the product (but also an object) of the planning process so that it provides richer information about the context of execution of the specific action. The context shall be particularly targeted towards mutual relation between the actions to be performed by individual actors and shall be used mainly for replanning and plan reparation purposes.

*The presented research has been supported by I-GLOBE (<http://i-globe.info/>), the US ARMY, CERDEC project (no.: N62558-06-P-0353). The project has been performed in cooperation with AIAI, University of Edinburgh.

Each individual actor - when agreeing with task implementation - adopts a commitment as a specific mental state, which represents the actor's proactive attitude towards goal implementation. The social commitment contains the following pieces of information:

- *commitment condition* that may be (i) a specific situation in the environment (such as completion of some precondition) or (ii) a time interval in which the action is to be implemented no matter what the status of the environment is or (iii) a combination of both.
- *decommitment conditions* specifying under which condition the actor is allowed to recommit from the commitment once the task is finished (e.g. notification) or once the task cannot be completed (e.g. a failure)

Michael Wooldridge in (Wooldridge, 2000) defines the commitments formally as follows:

$$\text{(Commit } A \psi \phi \lambda), \\ \lambda = \{(\rho_1, \gamma_1), (\rho_2, \gamma_2), \dots, (\rho_k, \gamma_k)\}, \quad (1)$$

where A denotes a committing actor, ψ is an activation condition, ϕ is a commitment goal, and λ is a convention. The convention is a set of tuples (ρ, γ) where ρ is a decommitment condition and γ is an inevitable outcome. The convention describes all

possible ways how the commitment can be dropped. Generally speaking, the actor A has to transform the world-state in such a way that the ϕ goal becomes true if ψ holds and any γ has not been made true yet. The actor is allowed to drop the commitment if and only if $\exists i : \rho_i$ which is valid. A decommitment is allowed provided that γ_i is made true.

Decommitment strategies represent the main distinction between a commitment and an action as a product of the multi-agent planning. The pivotal research problem when designing the multi-agent commitments-based planning algorithm is in designing agent's capability to reason not only about the actions, the conditions when the actions can be implemented and the quality of service but also about various decommitment rules and strategies to be adopted by the agent who commits to implementation of the particular action. More information about commitment-based planning in multi-agent systems and decommitment strategies can be found in (Komenda et al., 2008)

Another representation of commitments considering temporal account has been introduced in (Mallya et al., 2003). CTL (Emerson and Srinivasan, 1988) has been extended to capture features usually not considered in common approaches (but relevant for realistic environments), namely time intervals considered in commitments satisfaction, "maintenance" manner of commitments next to "achieve" manner of commitments and vague specification of time. The uncertainty in agent commitments has been studied in (Xuan and Lesser, 1999). The authors have extend the commitment by "... *uncertainty by explicitly describing the possibility of future modification/revocation of the commitment* ...". The focus has been put to the uncertainty in the quality of the commitment fulfilment (quality of service). Another aspects of commitments, such as start time, finish time, or duration has not been discussed.

In this article we will discuss the relaxation decommitment strategy based on Wooldridge's commitment representation, targeted specifically to the time interval in which the agent agrees to accomplish the commitment. We will discuss how these interval changes affect plan execution and how the potential changes of this interval can be represented in the commitment itself. The value of the use of social commitments in planning in dynamic, multi-actor environment has been documented on a series of empirical experiments presented in Section 3.

2 RELAXATION IN COMMITMENTS

The commitment time interval is usually captured by the commitment subject ϕ and specifies the booked time window for the commitment execution. The temporal uncertainty can be a part of the commitment subject definition (and thus the whole commitment has to be renegotiated in case of any change) or, more preferably, it can be included in the commitment as a special instance of a decommitment rule.

Let the commitment time interval $T_\phi = \langle t_s, t_e \rangle$, where the t_s is the starting time and the t_e is the ending time of the commitment time interval. The decommitment rules can then be described as:

$$\begin{aligned} &(\text{Commit } A \ \psi \ \phi \ \lambda), \\ &\lambda = \{ \{ ((t_s^{est} \neq t_s) \wedge (t_s^{est} \in T_s^{rlx}), t_s = t_s^{est}), \\ &\quad ((t_e^{est} \neq t_e) \wedge (t_e^{est} \in T_e^{rlx}), t_e = t_e^{est}) \}, \\ &\quad (t_s, t_e) \subset \phi, \end{aligned} \quad (2)$$

where t_s^{est} and t_e^{est} are the estimations of the real start and end of the activity and T_s^{rlx} and T_e^{rlx} are the agreed relaxation intervals for the start and end time.

The t_s , t_e , T_s^{rlx} and T_e^{rlx} are the parameters of the commitment negotiated and fixed in the planning (contracting) time. The t_s^{est} and t_e^{est} are the agent's estimates of the real start and end of the activity and can vary in time.

The proposed representation of temporal uncertainty in the commitment allows to allocate resources (by the means of agent commitments - one agent requests a resource that is provided by another agent) accordingly in uncertain environment. In the moment of allocation (e.g. negotiation of the commitment) the provider agent can make private estimation of the possible future progress and setup the relaxation intervals without the need to reserve redundant resources. The quality of estimation is the key issue and affects the quality and robustness of the overall plan. On the other hand, the requesting agent can select the provider according to the required flexibility or stability of the commitment.

2.1 Impact on Planning and Negotiation

Planning using agent commitments is an optimization method that takes into account individual goals and constraints of the actors. The plan is constructed by negotiation of agents, where the agents commit to perform specific actions under the agreed conditions. The key aspects of the negotiation are (i) establishing commitments that are expected to be kept (minimizing commitments violation) and (ii) minimize decommitment flexibility to avoid too much uncertainty.

Those two criteria are contradictory. In Section 2.2 we introduce a method that is focused on balancing the flexibility and the stability of the commitments.

During the commitment execution, the agent is able to update the estimations t_s^{est} and t_e^{est} for all future commitments based on the current (and past) performance and conditions. To enable propagation of this information the decommitment rules (Equation 2) can be extended with:

$$\lambda = \{ \dots, \begin{aligned} & ((t_s^{min} \neq \min(T_s^{relax})) \wedge (t_s^{min} \in T_s^{rlx}), \\ & T_s^{rlx} = \langle t_s^{min}, \max(T_s^{relax}) \rangle) \\ & ((t_s^{max} \neq \max(T_s^{relax})) \wedge (t_s^{max} \in T_s^{rlx}), \\ & T_s^{rlx} = \langle \min(T_s^{relax}), t_s^{max} \rangle) \\ & ((t_e^{min} \neq \min(T_e^{relax})) \wedge (t_e^{min} \in T_e^{rlx}), \\ & T_e^{rlx} = \langle t_e^{min}, \max(T_e^{relax}) \rangle) \\ & ((t_e^{max} \neq \max(T_e^{relax})) \wedge (t_e^{max} \in T_e^{rlx}), \\ & T_e^{rlx} = \langle \min(T_e^{relax}), t_e^{max} \rangle) \end{aligned} \quad (3)$$

In a complex interaction multi-agent scenario this information can help other agents to update their own estimations and potentially improve their plans.

Equation 3 ensures that the commitment keeps the previously agreed relaxation intervals but narrows them down according to the updated information. The relaxation intervals are updated to exclude parts where the committed action is unlikely to be executed.

2.2 Relaxation Interval

This section describes the method for setting the commitment relaxation interval inspired by the uncertainty handling in the PERT¹ diagrams. The PERT terminology operates with terms *optimistic time*, *most likely time* and *pessimistic time* that incorporates the uncertainty in the activity duration. Those terms are used to compute the properties for each activity in the project. In the network diagram representation, each activity is described by its duration, early start, early finish, late start, late finish and slack. Those parameters are very similar to the commitment representation of relaxation (the slack represents time reserve and is not used in commitments). The original PERT is usually used for a single project analysis, but setting of the commitment relaxation in a multi-agent system can be viewed from two points of view: (i) the requestor agent (RA) that maintains the commitments in the way similar to classical project management and (ii) the provider agent (PA) that maintains the commitments as independent entities agreed upon with (potentially) various requestors.

¹Program Evaluation and Review Technique

The role of the requestor agent is to ensure consistency between the dependent commitments constituted with particular provider agents. In this text we are focusing on the provider agent, whose role is to ensure commitments stability and flexibility with respect to the environment uncertainty and resources utilization. Those two points of view are very similar and the same strategies can be used for both.

The method of setting the relaxation interval parameters for i -th negotiated commitment is follows:

$$\begin{aligned} t_s(i) &= \max(t_e(i-1), t_s^{req}) \\ t_e(i) &= t_s(i) + t_d(i) \\ T_s^{rlx}(i) &= \langle \max(\min(T_e^{rlx}(i-1)), t_s^{req}), \\ & \quad \max(\max(T_e^{rlx}(i-1)), t_s^{req}) \rangle \\ T_e^{rlx}(i) &= \langle \min(T_s^{rlx}(i)) + t_d(i), \\ & \quad \max(T_s^{rlx}(i)) + t_d(i) + t_r * p_b^{worst}(i) \rangle, \end{aligned} \quad (4)$$

where $t_s^{req}(i)$ and $t_e^{req}(i)$ are the start and end times required by the RA, $t_d(i)$ is the nominal duration of the respective action, $p_b^{worst}(i)$ is the worst case estimate of probability of breakdown during the task execution and t_r is the reparation time of the resource.

The relaxation intervals defined by Equation 4 correspond to the classical PERT representation. In the case where the breakdown events happen more than once per one activity, the $p_b^{worst}(i)$ doesn't represent the probability value, but rather the relative number of breakdown events during the commitment execution.

3 EXPERIMENTS

The experiments evaluate three methods of handling uncertainty. The first method is based on statistical error evaluation and uses a constant safety margin without use of commitment relaxation rule. Other two are based on relaxation with different relaxation intervals estimation.

The experiments were performed on the scenario with one requestor agent (RA) and one provider agent (PA). The PA maintains a single resource that is used for tasks execution. The task execution is interruptible and only one task can be executed at any given time. The RA requests a set of 1000 tasks from the PA. The PA makes allocation for the tasks and proposes a commitment for each individual task. The commitment includes the start time, end time and relaxation intervals. The duration of the task is $t_d = 10$ seconds.

Environment uncertainty is modeled as a resource breakdown with a variable breakdown mean time t_b and reparation time of $t_r = 5$ seconds. An event sim-

ulation has been performed for 1000 randomly generated sets of breakdown events for each experiment setting. The commitment parameters were computed as follows:

(M₁): Constant – this method extends the duration of the task by the relative reparation time computed by the probability of the breakdown for each task. The start time of the commitment $t_s(i) = t_e(i - 1)$ and the $t_e(i) = t_s(i) + t_d(i) + t_r * p_b(i)$, where $p_b(i)$ is the probability of the breakdown during the commitment execution.

(M₂): Linear – this method doesn't change the duration of the activity, so the $t_s(i) = t_e(i - 1)$ of the previous commitment and the $t_e(i) = t_s(i) + t_d(i)$. The relaxation interval for the decommitment rule is computed as described in Section 2.2.

Commitment parameters computation are based on the estimation of breakdown mean time $\bar{e}_{est} = 15$ and known reparation time $t_r = 5$ seconds. The p_b^{worst} estimation is set to \bar{e}_{est} . We measure the robustness of the commitments and the resource utilization (total execution time) under various conditions generated by several environment uncertainty models.

The uncertainty models have been generated using three methods. Each method produces an event sets with various properties of mean value \bar{e} and standard deviation σ . The environment uncertainty models are the following:

(U₁): Deterministic – breakdown events are generated evenly with t_b period. This method produces constant $\bar{e} = t_b$ and $\sigma = 0$.

(U₂): Gauss – generates a set with normal distribution with $\bar{e} = t_b$, $\sigma = t_b/10$ with a delay between two subsequent events restricted to $\langle 0, 2 * t_b \rangle$.

(U₃): Uniform – generates a set of uniformly distributed events with $\bar{e} = t_b$ with a delay between two subsequent events restricted to $\langle 0, 2 * t_b \rangle$.

The experiments were run with \bar{e} value varied according to Table 1. The random sets were generated uniformly in this interval to evaluate the robustness of the commitment relaxation setting methods. Because of the relatively small event sets, the generated pseudo-random values don't fit exactly to the desired parameters (especially mean time value). The real \bar{e} of each set has been computed within the simulation run and corresponds to the x-axis in the provided evaluation figures.

3.1 Observations

The figures below show the number of violated commitments, tardiness of the commitments, and total execution time for all methods under various uncertainty

Table 1: Properties of the breakdown distributions.

distribution	$\bar{e}(s)$	$\sigma(s)$
U_1	10 ~ 20	0
U_2	10 ~ 20	1.5
U_3	10 ~ 20	8.66

settings. The total execution time of all commitments is $\sum t_d(i) = 1500$ seconds and represents the ideal execution duration in breakdown-free environment. The length of the plan represents the end time of the last commitment for the M_1 method and the latest time of the relaxation intervals for the M_2 method. The plan length is influenced by the experiment settings and in our case it lengthens the plan for both methods by 50% (caused by the parameters \bar{e} , t_d and t_r).

During the simulation, the agents kept all the commitments as agreed at the beginning. The experiments show the impact of non-accurate estimation on the error mean time. As expected, both methods provide good results when $\bar{e} > \bar{e}_{est}$. When the breakdown mean time is shorter, both methods start to generate commitment violations.

Figure 2 shows the mean time influence on the number of violations. The commitment is violated when it cannot be finished within the agreed limits (t_e for M_1 and relaxation interval T_e^{rlx} for M_2). The robustness of the method M_1 is very limited. It provides good results only for deterministic uncertainty U_1 with $\bar{e} \geq \bar{e}_{est}$. When $\bar{e} < \bar{e}_{est}$, all the commitments are violated (Figure 2–a). The M_2 method provides better results in the left-hand part of the graph ($\bar{e} < \bar{e}_{est}$) because of greater safety margin (caused by p_b^{worst}) but still converges quickly to the 100% violated commitments. For normal uncertainty U_2 the situation changes. The M_1 method fails in the entire range of \bar{e} and there is a low amount of non-violated commitments even in the range $\bar{e} > \bar{e}_{est}$ (see Figure 2–b). The M_2 method provides a minimal amount of violated commitments in the region $\bar{e} > \bar{e}_{est}$ and $\bar{e} \sim \bar{e}_{est}$ and the number of violated commitments slowly grows in the range $\bar{e} < \bar{e}_{est}$. In the case of uniform uncertainty U_3 both methods fail (Figure 2–c). The lowest number of violated commitments is in the right-hand region and it goes from approximately 50% to more than 80% for $\bar{e} = \bar{e}_{est}$. The method M_2 provides good results for $\bar{e} \geq \bar{e}_{est}$. The number of violations grows with descending \bar{e} . For $\bar{e} = \bar{e}_{est}$ the average number of violated commitments is about 50%.

The total execution time is presented in Figure 1. Commitment execution is not started before the agreed time (t_s for M_1 and the relaxation interval T_s^{rlx} for M_2), so the execution time mainly corresponds to the number of violated commitments. In case of the




Figure 1: Total execution time with (a) deterministic mean time of breakdowns, (b) normally distributed mean time of breakdowns ($\bar{e} = t_b$, $\sigma = t_b/10$), (c) uniformly distributed mean time of breakdowns.




Figure 2: Number of violations with (a) deterministic mean time of breakdowns, (b) normally distributed mean time of breakdowns ($\bar{e} = t_b$, $\sigma = t_b/10$), (c) uniformly distributed mean time of breakdowns.

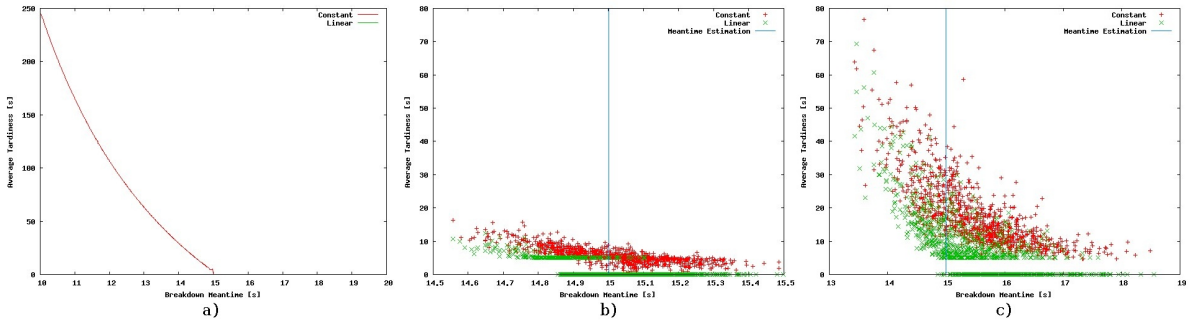


Figure 3: Average tardiness with (a) deterministic mean time of breakdowns, (b) normally distributed mean time of breakdowns ($\bar{e} = t_b$, $\sigma = t_b/10$), (c) uniformly distributed mean time of breakdowns.

M_1 method, the minimal execution time is equal to the plan length (1500). The M_2 method execution time converges to the $\sum t_d(i)$ for the increasing \bar{e} .

For deterministic uncertainty U_1 both methods provide hyperbolic growth of the total execution time with decreasing \bar{e} (see Figure 1–a). The relatively small difference between the methods is caused by the fast growth of the violated commitments of method M_2 and the low tardiness of the M_1 commitments in the range where M_2 keeps the number of violated commitments low. For normal uncertainty U_2 the total execution time grows almost linearly with decreasing \bar{e} . The difference between the methods is given by the

difference in the number of violated commitments, which is considerably higher for M_1 . For \bar{e} smaller than the depicted value range, the execution time converges to the execution time curve of the U_1 (as the number of violations grows). The same situation occurs for the case of uniform uncertainty U_3 . The disruption of the execution time curve of method M_1 is given by high variation of the number of violated commitments. The total execution time of the M_2 is similar to the total execution time for the other two environment settings. The only difference is given by small variation ($\pm 2\%$ or less) of the execution time for the particular \bar{e} . This variance depends on the vari-

ation of number of violated commitments for this \bar{e} across the simulation runs.

The average tardiness of the commitment completion presented on Figure 3 is computed for all violated commitments. The non-violated commitments are not taken into account. If there is no violated commitment, the average tardiness is set to zero. For deterministic uncertainty U_1 the results of both methods are very similar. The tardiness grows with decreasing \bar{e} in a similar way as the total execution time. For normal uncertainty U_2 both methods provide low tardiness that again converges to the curve for the U_1 environment setting for small values of \bar{e} (the convergence is not captured by the Figure 3–b). Similarly to the total execution length, the disruption of the curves is given by the variation of the number of violated commitments. For uniform uncertainty U_3 the average tardiness grows faster with the decreasing \bar{e} (see Figure 3–c). The method M_1 provides relatively high average tardiness of the commitments even in the region $\bar{e} > \bar{e}_{est}$ and $\bar{e} \sim \bar{e}_{est}$. The method M_2 provides better results and the tardiness grows mainly in the range $\bar{e} < \bar{e}_{est}$.

4 CONCLUSION

We have presented the social commitment representation for multi-agent planning and plan execution in the distributed domain with environment featuring uncertainty. We have defined a relaxation decommitment strategy targeted specifically to the time interval in which the agent agrees to accomplish the commitment. The relaxation strategy setting has been experimentally evaluated and compared with a basic method utilizing fixed commitments. The experiments have proved that incorporating potential relaxation brings certain benefits in comparison to the constantly evaluated safety margins.

The basic method M_1 is suitable mainly for deterministic environment U_1 where the relaxation decommitment strategy method M_2 brings no significant improvement. Extending the safety margins in both methods can scale the results towards lower \bar{e} but lengthen the plans (and also the total execution time for M_1). For environments U_2 and U_3 , increasing the safety margin brings no significant advantage because of higher distortion of the breakdown distribution.

From the point of view of the number of violated commitments, which is extremely important in the multi-actors scenarios, the method M_1 fails for U_2 and provides even worse results for U_3 . In this case, the relaxation decommitment method M_2 is beneficial for U_2 and keeps certain advantages even in U_3 , where

the average number of violated commitments is about 50%.

Another advantage of the M_2 method is its robustness. We have experimentally proved that the total execution duration and commitment tardiness does not depend very much on the breakdown distribution function (the results of experiments don't differ by more than 2%). With the increasing \bar{e} the total execution time converges to $\sum t_d(i)$, which is the minimal possible execution time. Due to the start time and end time relaxation ability, the method enables both optimistic and pessimistic execution without breaking the commitments. The relaxation decommitment strategy greatly increases the flexibility, stability and robustness of the agents' social commitments in the dynamic uncertain environment.

REFERENCES

- Emerson, E. and Srinivasan, J. (1988). Branching time temporal logic. In *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, School/Workshop*, volume 354 of *LNAI*, pages 123–172. Springer-Verlag.
- Komenda, A., Pechoucek, M., Biba, J., and Vokrinek, J. (2008). Planning and re-planning in multi-actors scenarios by means of social commitments. In *Proceedings of the International Multiconference on Computer Science and Information Technology (IMCSIT/ABC 2008)*, volume 3, pages 39–45. IEEE.
- Levesque, H., Cohen, P., and Nunes, J. (1990). On acting together. In *AAAI-90 Proceedings, Eighth National Conference on Artificial Intelligence*, volume 2, pages 94–99, Cambridge, MA, USA. MIT Press.
- Mallya, A., Yolum, P., and Singh, M. (2003). Resolving commitments among autonomous agents. In Dignum, F., editor, *Advances in Agent Communication Languages, ACL 2003*, volume 2922 of *LNCS*, pages 166–82. Springer-Verlag, Berlin, Germany.
- Wooldridge, M. (2000). *Reasoning about Rational Agents*. Intelligent robotics and autonomous agents. The MIT Press.
- Wooldridge, M. and Jennings, N. (1999). Cooperative problem solving. *Journal of Logics and Computation*, 9(4):563–594.
- Xuan, P. and Lesser, V. R. (1999). Incorporating uncertainty in agent commitments. In *In: Proc. of ATAL-99*, pages 221–234. Springer-Verlag.