Österreichisches Forschungsinstitut für / Austrian Research Institute for / Artificial Intelligence

TR–2004–03

Bernhard Jung, Paolo Petta

Improving upon the TAEMS/DTC framework in the context of coordinated scheduling

• Freyung 6/6 • A-1010 Vienna • Austria •

• Phone: +43-1-5336112 •

• mailto:sec@oefai.at •

• http://www.oefai.at/oefai/ •

Österreichisches Forschungsinstitut für / Austrian Research Institute for / Artificial Intelligence

TR–2004–03

Bernhard Jung, Paolo Petta

Improving upon the TAEMS/DTC framework in the context of coordinated scheduling

The Austrian Research Institute for Artificial Intelligence is supported by the Federal Ministry of Education, Science and Culture.

Citation: Jung B., Petta P.: Improving upon the TAEMS/DTC framework in the context of coordinated scheduling. Technical Report, Österreichisches Forschungsinstitut für Artificial Intelligence, Wien, TR-2004-03, 2004. to appear in Proc. of AT2AI-4: Fourth International Symposium "From Agent Theory to Agent Implementation".

Improving upon the TÆMS/DTC framework in the context of coordinated scheduling

Bernhard Jung and **Paolo Petta**
Austrian Research Institute for Artificial Intelligence (ÖFAI)*
Freyung 6/6, A-1010 Vienna, Austria
email: {bernhard.jung,paolo}@oefai.at

Abstract

TÆMS, DTC and GPGP constitute an evolved framework for coordination in multi-agent systems. In this paper, we focus on inconsistencies and semantic interpretation problems encountered during implementation of a criteria-driven scheduler. We try to disambiguate and simplify concepts and propose extensions and new features for TÆMS and DTC to improve understanding, use, and integration of the framework in an agent architecture. We propose how to modularise TÆMS/DTC to form a kind of construction kit for local agent coordination and control and extend the scope to also cover domain-dependent context. Throughout, we aim to simplify application and integration of the TÆMS/DTC framework while preserving its core ideas. We also locate our work in the context of current efforts to develop a unified view of coordination in MAS.

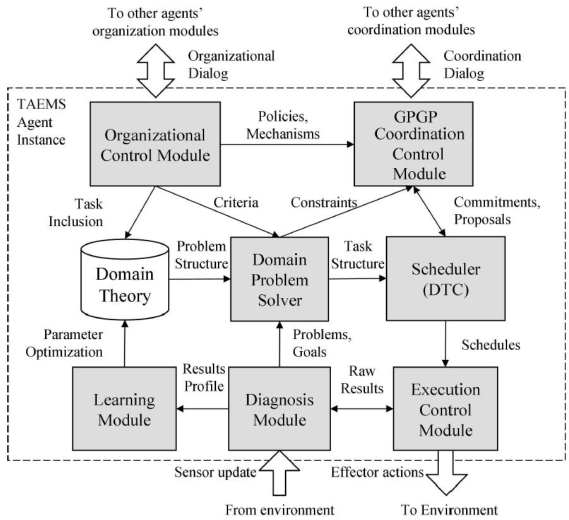

Figure 1: A TÆMS-based Agent Architecture (taken from [Wagner et al. 2003]).

1 Introduction

Three main aspects of the coordination problem in multi-agent systems regard: *Coordination between multiple tasks and goals within a single agent instance:* agents pursuing multiple tasks in order to achieve one or more goals have to *sequence their actions* in respect to ordering constraints, task interdependencies, time limits and goal priorities; *Coordination between multiple agents:* multiple agents interacting in a common environment have interests to coordinate their actions to achieve shared goals, minimise redundant efforts, or avoid obstructive actions; and *Coordination of resource usage:* when agents interact with their environment, the modeling of resources and resource usage becomes important and thus coordination of resource usage, either between agents or within a single multi-tasking agent.

Among the technologies addressing these issues, the TÆMS (Task Analysis, Environment Modeling and Simulation)[1] framework, taken jointly with GPGP (Generalized Partial Global Planning)[2] and DTC (Design-to-Criteria Scheduling)[3], forms an approach that comes to lie between design-time solutions—e.g., organisational restrictions or problem specific coordination protocols—that do not (fully) address problems and opportunities arising dynamically, and run-time techniques that have to struggle with the huge degrees of freedom of interaction space and tend to introduce significant communication overhead, while still allowing for short-term planning only.

TÆMS provides a framework to model hierarchical task-structures for multiple agents. Beyond simply structured decomposed tasks and goals, it supports descriptions of interdependencies between subtasks (within the same task or across different tasks; within one or between different agents) and of interrelationships between actions and resources. TÆMS uses a quantitative approach to describe tasks and methods in the three dimensions of quality, cost and duration. Uncertainties are taken into account by using discrete probability distributions. GPGP

*The Austrian Research Institute for Artificial Intelligence is supported by the Austrian Federal Ministry for Education, Science and Culture and by the Austrian Ministry for Transport, Innovation and Technology.
Part of this research was carried out in the context of the research project S9106-N04 of the FWF Austrian Science Fund.

[1] http://dis.cs.umass.edu/research/taems/
[2] http://mas.cs.umass.edu/research/gpgp/
[3] http://mas.cs.umass.edu/research/dtc/

uses TÆMS in multi-agent planning to identify and generalise types of coordination relationships. DTC exploits TÆMS to find particular sequences of actions to achieve agents' goals that respect shared actions, time and resource restrictions. A third area of TÆMS applications is simulation of multi-agent systems (MASS – The Multi Agent Survivability Simulator)[4]. TÆMS, DTC and GPGP are typically deployed together with a domain planner, an execution and monitoring component, and some additional modules (figure 1, see also [Horling et al. 2002]). This technology has been successfully applied in several implementations and domains including emergency response control [Phelps et al. 2003]; supply chain management [Wagner et al. 2002]; information gathering [Lesser et al. 2000]; domotics [Lesser et al. 1999]; and distributed sensing [Horling et al. 2001].

Even so, there exist several semantic ambiguities and inconsistencies of use of TÆMS concepts across TÆMS itself and within its use in the DTC framework. In part these came about as a consequence of steady extension of the framework, and to some extent they are documented and acknowledged explicitly (e.g. [Horling et al. 2003, Wagner 2003]); but some further issues arose during our work on the implementation of a DTC scheduler. The following section describes TÆMS and DTC briefly; afterwards, we address some of these problems encountered and propose our extensions and improvements towards a modular and flexible DTC architecture.

2 TÆMS & DTC

The description of the TÆMS modelling language is consolidated in the TÆMS White Paper document [Horling et al. 2003], supplemented by [Wagner 2003] for DTC specifics. TÆMS shows the following characteristics: *Hierarchical task-structures* define problems in terms of goals, tasks and sub-tasks and their task-subtask relationships; *Quality accumulation functions (QAFs)* describe how qualities of tasks result from those of subtasks (possibly imposing ordering restrictions); *Resources* model agents' interactions with their environment, where actions can produce or consume resources (and out of bounds resources can affect actions in a "soft" way); *Interrelationships (IRs)* describe hard (enabling/disabling) or soft (negative/positive) effects of task executions on other tasks or resources; *Quantitative and uncertain descriptions* of tasks and interrelationships along the dimensions of quality, cost, and duration, by means of discrete probability distributions[5]; *Domain-independence* achieved by abstract, generic concepts to describe tasks, sub-tasks, and interrelationships and by using simple execution profiles with quality, cost and duration for actions; *Multi-agent capability* by modeling task-structures for several agents and describing interrelationships and commitments between agents; *Plaintext and simple syntax* for transparent, platform-independent representation.

TÆMS task structures can be regarded as descriptions of coordination, planning and scheduling problems to be solved by DTC, if partially, as it only considers the commitments coordination restrictions when generating multi-agent schedules. DTC evolved from Design-to-Time [Garvey & Lesser 1995] and is described in several publications (including [Wagner et al. 1998a, Wagner et al. 1998b, Wagner 2000]). Throughout DTC, criteria are used to specify relative solution preferences according to quality, cost, duration, and uncertainties along these dimensions. Additional reward can be assigned upon respecting cost/duration limits or some minimum quality. The slider metaphor is often used to describe these criteria (figure 2). The basic DTC algorithm comprises alternative generation (planning), method ordering (scheduling), and application of critics (schedule improvement): due to the complexity of planning and scheduling, an iterative approximation approach driven by heuristics is used to achieve anytime behaviour.

[4] http://mas.cs.umass.edu/research/mass/

[5] The use of probability distributions is not *"perfect science"* [Wagner et al. 1998b].

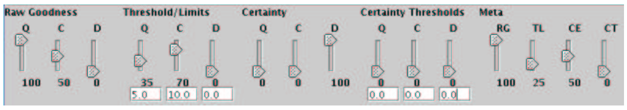

Figure 2: The slider metaphor to describe criteria (adapted from [Wagner et al. 1998b]).

3 Inconsistencies and semantic issues

As stated before, inconsistencies and semantic interpretation problems concerning scheduling currently exist in TÆMS, with even some of the documented ones lacking a detailed discussion. In this section, we review problems we encountered during the implementation of a DTC scheduler. Our analysis also considers empirical results obtained with the DTC implementation of the Java Agent Framework[6] (JAF-DTC). [7].

3.1 Quality accumulation, task ordering & activation

QAFs play a central role in the capability of TÆMS to describe alternative ways of achieving a task. As stated in [Wagner et al. 2003], QAFs evolved from plain quality accumulation to complex functions that additionally restrict subtasks ordering, so as to simplify modelling. A question not fully addressed in the White Paper is the *time* of quality accumulation, needed to activate interrelationships. For *sequencing* and some other QAFs, accumulation time is definitely after execution of the last sub-task, but for others (e.g., *sum*) quality accumulation can happen either *once* after the execution of any subset of sub-tasks, or can change *continuously* with every executed sub-task[8]. A task with *once* semantics can be seen as requiring an explicit aggregation of the subtask results[9]; a *continuous* task reflects more of a kind of "world state" view.

We suggest that a clear separation of the two aspects, quality accumulation and ordering restrictions, together with the definition of the accumulation time, can clarify the

[6] http://mas.cs.umass.edu/research/jaf/

[7] For detailed results, see [Jung 2003].

[8] The JAF-DTC implementation does not handle the activation of IRs by partially executed tasks in a consistent way.

[9] Which can be modelled by introducing an additional method of negligible duration.

semantics of QAFs and allow straightforward introduction of new *complex QAFs* as symbol triplets covering the following aspects: **Quality calculation**, a function deriving a task's quality out of those of its sub-tasks; apart from *sum*, *min*, and *max*, more complex (predefined)functions can be used. **Ordering restrictions** express explicitly the restrictions of *seq* QAFs; apart from simple sequencing, partial ordering or even more complex orderings can be covered. Ordering restrictions verify ordering of partial schedules and assess whether appending a certain method is admissible. Using complex ordering restrictions, a specific subtask can be designated as "aggregation task" to be executed last, independently from the ordering of the other subtasks. The **Accumulation type** defines the accumulation time of the task as *once* or *continuous*. For *once*, accumulation has to be triggered explicitly when the task's quality is needed, either to contribute to a super-task or to activate interrelationships[10].

3.2 Soft IRs and their effects

Soft IRs affect quality, cost and duration according to $v_{new} = v \pm v * p * sq/smq$. v is the value of quality, cost, or duration; p is the power factor of the IR, and sq/smq the source factor scaling the effects by the relative quality (actual vs. maximum possible quality) of the source of the activating task.

The formula raises several issues: **Negative values** might occur due to power or source factors > 1. To avoid them, both factors can be restricted to ≤ 1, or the range of v_{new} can be limited to non-negative values. The main difference between these two solutions lies in the steepness of the effects for small source factors. For **source factor calculation** the maximum achievable quality has to be known, either considering all possible facilitating IRs or just the maximum of the base distribution. The latter can lead to actual source factors > 1. To avoid such problems, the source factor can be redefined as $min(1, sq/smq)$. The TÆMS White Paper specifies the calculation of **multiple effects** to occur in the order of the respective IR activation times. However, a closer look at the formula reveals that the calculation is commutative and the sequence of the calculations thus does not matter. Even so, for multiple IRs pointing to the same node it has to be clarified whether the accumulated value v_{new} is to be the product or the sum of the contributions of the individual IRs[11].

Experiments with JAF-DTC show that the problems with negative values are not addressed; source factor scaling is only implemented for simple unaffected methods (for tasks and affected methods a constant value of 1 is taken), and that multiple effects use the already affected value for calculation.

3.3 Interrelationships pointing to tasks

[Garvey & Lesser 1995] define the semantics of IRs pointing to tasks as: "A relationship from Task A to Task B is translated to relationships from Task A to all methods below Task B". However, direct effects on task qualities can be useful in certain models (while cost and duration always result from subtasks), where soft direct interrelationships affect the calculated quality and hard direct IRs enable or disable the accumulation process itself.

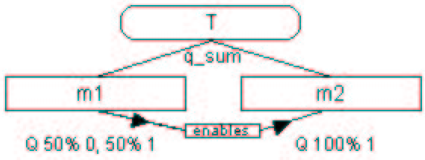

Figure 3: A simple TÆMS structure with dependent subtasks.

3.4 Resource usage

TÆMS supports the *per time unit* and *duration independent* resource models. While the latter is consistent with atomic actions, the *per time unit* model is not, as resource usage by methods and the effects of resources on methods depend on the method's quality and the resource level at each time tick. For a correct simulation of the *per time unit* model, TÆMS would need a complex method model allowing to describe method's execution profiles and the impact of resources on methods on a time tick basis. Our work currently ignores this model due to its impact on "correct" simulations[12].

3.5 Uncertainties and the independence assumption

The introduction of uncertainty in TÆMS had a strong impact on the handling of TÆMS structures in the context of scheduling ([Garvey & Lesser 1995, Wagner & Garvey 1997]). The maintenance of probability distributions can become difficult and requires substantial memory space for combinations of even small distributions. A compaction of distributions is used in DTC scheduling to cope with this problem, at the cost of losing precision [Wagner 2000][13]. DTC also has to handle distributions for decisions (e.g., whether an IR is activated, a resource is out-of-bounds, a deadline is missed), during the calculation of IR effects (using uncertain qualities, power and source factors, resource amounts, and IR activation times and probabilities), etc.

The basic assumption of independence of quality, cost, and duration distributions may not be accurate for all kinds of actions, but simplifies working with distributions. However, TÆMS structures can introduce explicit dependencies. Figure 3 shows a simple example with a task T with QAF *sum* and two methods, $m1 and m2$, that are connected by an *enables IR*. The quality of $m1$ is (50% 0, 50% 1), thus enabling $m2$ in 50% of all cases and the quality of $m2$

[10] In DTC, the time of accumulation can be derived from the methods in the alternative actually scheduled.

[11] I.e., either $v_{new} = v \pm v * p_1 * sq_1/smq_1 \pm v * p_2 * sq_2/smq_2 \pm v * p_3 * sq_3/smq_3 \pm \ldots$
or $v_{new} = v * (1 \pm p_1 * sq_1/smq_1) * (1 \pm p_2 * sq_2/smq_2) * (1 \pm p_3 * sq_3/smq_3) * \ldots$

[12] The JAF-DTC implementation supports both models, but reduces probability distributions of durations and amounts to expected values and thus simplifies the decisions for out-of-bound resources. The calculation of limiting IRs is faulty at times and can lead to negative costs and durations.

[13] The compaction algorithm used by the JAF-DTC implementation is undocumented.

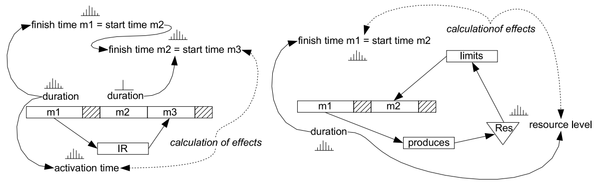

Figure 4: Dependencies for delayed IRs and resource levels.

is (50% 0, 50% 1). Ignoring the dependency between the two sub-qualities, overall quality results as (25% 0, 50% 1, 25% 2), while the result considering the dependency is (50% 0, 50% 2), as zero quality of *m*1 does not enable *m*2 and maximum quality is accumulated otherwise.

Apart from the sub-task dependency of tasks, we will point out two other examples (Figure 4), where the independence assumption does not hold. If an *IR is delayed* and its applicability to a certain method has to be tested, activation and start times of the method have to be assessed: Both variables can be distributions and they are not independent, as both depend on the end time of the activating method. If a resource was changed by a method with uncertain duration according to the *per time unit* model, the *resource level* depends on the method's duration. If a subsequent method can be limited by this resource, the effect depends on its start time. An early start time due to short duration of the first method implies little change of the resource level, while a later start time entails a larger change of the resource level. Dependencies thus do not just occur along a single dimension, as quality can become dependent on time.

The JAF-DTC implementation addresses the independence assumption by using a tree-based in-context analysis to correct the qualities *after* schedule construction (Section 4.2.3 in [Wagner 2000]). This analysis seems to cover only the problem of sub-task dependencies. The other two problems do not occur in JAF-DTC because delay distributions (for delayed IRs) and amount and duration distributions (for resource level calculation) are reduced to their expected values.

4 Improvements on TÆMS and DTC

The analysis of TÆMS and DTC has produced a number of ideas how to improve and extend semantics, that are being integrated into our implementation of a criteria-driven scheduler [Jung 2003]. Additionally, the following desiderata are known from the literature:

[Wagner et al. 2001b] states that TÆMS lacks a meta language to allow to represent "complex, situation-dependent relationships between functional decompositions of agent actions". A meta language can also be used to describe the semantics of QAFs, IRs, resource models, etc. The generation of alternatives using a best-*N* pruning leads to bad results in certain, highly restricted task-structures. A sampling of alternatives can be a way to overcome this problem ([Wagner & Lesser 1999]). [Raja & Lesser 2001] propose DTC extensions promoting scheduling effort, slack, and horizon to first class objects, to support meta-level control. In this context, a modular and tunable scheduler can be useful, as different environments require or allow for special reasoning. E.g., in highly uncertain environments, complex resource simulation is counterproductive, as schedules are likely to fail soon and simulation time may be better allocated to execution.

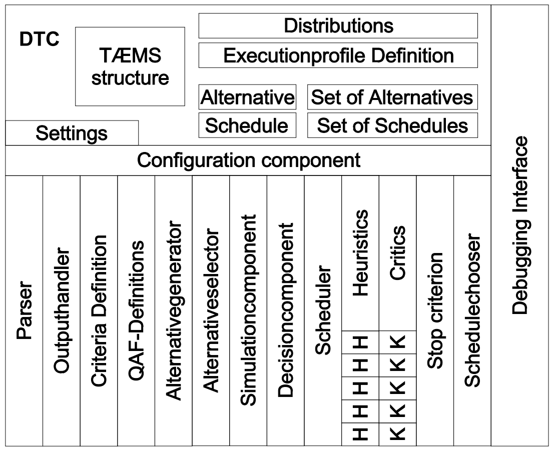

Figure 5: The VIE-CDS architecture

These observations substantiated our plans to realise a modular architecture for criteria-driven scheduling, VIE-CDS. The following gives an overview of the concepts integrated into the implementation.

Our modular extensible implementation uses a small TÆMS/DTC core and many exchangeable plug-ins that encapsulate input/output-handlers, strategies, calculations and selectors (figure 5). This architecture allows to easily integrate new components, such as strategies for alternative generation, new output formats, resource simulations of differing complexity, a partial order scheduler instead of heuristic scheduling, new heuristics, special purpose critics, etc. Dynamic reconfiguration is supported by the configuration component.

Execution profiles can be defined in a flexible way. The dimensions of quality, cost, and duration can be reduced or extended, if with some limitations: The dimension duration is mandatory, other dimensions being categorised as primary (handled by QAFs) ad secondary (handled as costs and are combined by adding up). Flexible profiles require flexible criteria definitions allowing to cover new dimensions. In addition to "slider criteria", VIE-CDS allows plug-in integration of other utility functions.

Complex QAFs are used as described in section 3.1. Quality functions and ordering restrictions can be extended easily by providing additional classes and registering with the configuration component.

The calculation of soft IR effects has been harmonised for methods; tasks can be parameterised to support different calculations as described in section 3.2.

Debugging support and a graphical user interface help users and researchers to analyse the behaviour of the DTC and its interplay with the other components of VIE-CDS.

The integration of these extensions and the new architecture preserve the core of DTC. The current implementation of VIE-CDS is about an order of magnitude slower than JAF-DTC, due to the usage of Java (vs. C/C++) and modularity overheads, with the parser contributing substantially for smaller TÆMS structures. While VIE-CDS does not yet provide an explicit meta language, but already supports additional semantics e.g., for QAFs integrated via plug-ins.

5 Conciliation with related work

The problem of coordination in MAS has been recently surveyed, reviewed and organised by members of the AgentLink working group on "communication, coordination and collaboration" of information agents: [Schumacher 2001, Omicini & Ossowski 2003] focused on objective vs. subjective views on coordination, and [Omicini et al. 2004] propose to recast coordination based on Activity Theory (AT) and to integrate objective coordination into the subjective view of FIPA infrastructures (see also e.g. [Ricci et al. 2004, Viroli et al. 2004]). An excellent deep analysis of coordination issues that can be addressed using the TÆMS framework, is given by [Lesser 1998]. Within the context of subjective vs. objective coordination and AT, which analyses every collaborative activity working with coordination artefacts at the three hierarchical levels of *co-construction*, *co-operation* and *co-ordination*, we can rethink the role of TÆMS as follows:

TÆMS task-structures, along with organisational knowledge and (partial) schedules, can be regarded as coordination artefacts. At the level of *co-construction*, they serve the purpose of erecting a (possibly partial) global view (by selection from a library or assembly of pieces provided by individual agents) and analysing inter-agent dependencies. This analysis then helps at the level of *co-operation* to form commitments or refine task-structures to avoid potentially hindering or hindered tasks. Such constraints lead to a coordinated schedule, that—at the level of *co-ordination*—is executed and monitored.

Reflection on the coordination artefacts (task-structures, organisational knowledge, schedules) wrt. task performance characteristics, unknown coordination relations, and characteristic dynamic properties, can lead to *routinisation* by adaptation of the TÆMS plan library and allowed coordination strategies pursued by the scheduler at *co-construction* level, and refinement of non-local effects, qualities and schedule effort (in terms of scheduling strategies and heuristics) at the *co-operation* level. In particular, VIE-CDS provides support for refinement at the *co-operation* level by virtue of its modular architecture.

Throughout, the TÆMS framework so far played a subjective role in coordination. Objective coordination can be envisioned by using a "scheduler-as-a-service" concept, that independently from the agent itself provides schedules and integrates overall system information unknown to (or unrecognised by) a single agent, e.g. reduces schedule effort when rescheduling happens too often. As pointed out in [Lesser 1998], the design space for such functionality in fact encompasses individual agents, agent task groups (cf. also the TeamCore architecture assigning dedicated team agents for this purpose [Tambe & Zhang 2000]), and more general monitoring services, each capable to contribute differently according to the different perspectives.

6 Summary & future work

We have proposed clarifications and extensions to the TÆMS framework, developed a new modular and flexible architecture, and integrated new features. With the completed implementation, we are now moving on to collect practical experiences with the integration of VIE-CDS into MAS. Main entries on our current to-do list are summarised below.

Integration in agent middleware: To boost the usage of VIE-CDS and to overcome limitations of the current round-robin scheduler, we plan to integrate VIE-CDS with JADE[14]. Feedback from the large JADE user community shall help to evaluate the usefulness of our architecture and components and identify requirements for further work.

Negotiation interface: The initial idea for negotiation support was already laid out in [Garvey et al. 1994]. A negotiation interface between the scheduler and a domain planner seems to be necessary for good integration, nevertheless it seems to be hard to achieve. Schedule annotations as done by JAF-DTC may not be sufficient; possibilities to control and drive the scheduling process into a certain direction might be also necessary. Such simple annotations or dynamic deadline relaxations could be a starting point to investigate the requirements and usefulness of a negotiation interface.

Domain-dependent contextual decisions: The only context that TÆMS can access from within a given TÆMS structure is the set of already executed methods and the implicit context that is inherent in the task-structure itself. Accessing external context to decide on the exact behaviour of certain QAFs or interrelationships could make task structures more general and more expressive, by taking on part of the responsibility that currently lies with the domain planner. The idea would be to declare extension points for domain- and context-dependent user-defined QAFs, ordering constraint functions, and interrelationships. The integration of such objects would not affect the DTC core and allow to choose a specific behaviour depending on an external context.

Modeling support: To exploit complex QAFs and specifically complex ordering constraints, modeling support can be useful. This support could utilise external context for deciding on "practices" for tasks. Contextual graphs ([Brezillon 2003]) could be a starting point for such modeling.

Execution support: TÆMS/DTC provides no execution and monitoring support. Some related improvements have already been presented in [Wagner et al. 2001b], such as extending schedule descriptions with domain-dependent data.

In addition, we are evaluating work on contingency analysis and meta-level control as well as TÆMS extensions such as parallel execution and iterations for future integration. Finally, we anticipate that the line of analysis presented in section 5 will evolve into a main driver of future activities.

[14] http://sharon.cselt.it/projects/jade/

References

[Brezillon 2003] P. Brézillon: Context dynamic and explanation in contextual graphs. in P. Blackburn, C. Ghidini, R.M. Turner, F. Giunchiglia (eds.): *Modeling and Using Context (CONTEXT-03)*. LNAI 2680, Springer-Verlag Heidelberg Berlin. pp. 94–106. 2003.

[Garvey et al. 1994] A. Garvey, K. Decker, V. Lesser: A Negotiation-based Interface Between a Real-time Scheduler and a Decision-Maker. UMASS Dept. of CS TR-1994-008. 1994.

[Garvey & Lesser 1995] A. Garvey, V. Lesser: Design-to-time Scheduling with Uncertainty. UMASS Dept. of CS TR-1995-003. 1995.

[Horling et al. 2001] B. Horling, R. Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins, V. Lesser: Distributed Sensor Network for Real Time Tracking. *Proc. of the 5th Intl. Conf. on Autonomous Agents*. ACM Press. pp. 417–424. 2001.

[Horling et al. 2002] B. Horling, V. Lesser, R. Vincent, T. Wagner: The Soft Real-Time Agent Control Architecture. UMASS Dept. of CS TR-2002-014. 2002.

[Horling et al. 2003] B. Horling, V. Lesser, R. Vincent, T. Wagner, A. Raja, S. Zhang, K. Decker, A. Garvey: The Taems White Paper. 16.04.2003. http://dis.cs.umass.edu/research/taems/white/.

[Jung 2003] B. Jung: VIE-CDS: a modular architecture for criteria-driven scheduling, diploma thesis, Vienna University of Technology, 2003 (in german).

[Lesser 1998] V. Lesser: Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture. *Autonomous Agents and Multi-Agent System*. **1(1)**:89–111, Kluwer Academic Publishers, 1998.

[Lesser et al. 1999] V. Lesser, M. Atighetchi, B. Benyo, B. Horling, A. Raja, R. Vincent, T. Wagner, P. Xuan, S. Zhang: The UMASS Intelligent Home Project. Etzioni O., Müller J.P., Bradshaw J.M.(eds.): *Proc. of the 3rd Intl. Conf. on Autonomous Agents (Agents-99)*, pp. 291–298, May, 1999.

[Lesser et al. 2000] V. Lesser, B. Horling, F. Klassner, A. Raja, T. Wagner, S. Zhang: BIG: A Resource-Bounded Information Gathering and Decision Support Agent. *Journal of Artificial Intelligence*, **118**:1–2, pp. 197–244, May 2000.

[Omicini & Ossowski 2003] A. Omicini, S. Ossowski: Objective versus Subjective Coordination in the Engineering of Agent Systems. in Klusch M. et al.(eds.): *Intelligent Information Agents: The AgentLink Perspective*. LNAI 2586 (State-of-the-Art Survey). Springer-Verlag Berlin Heidelberg, pp. 179–202, March 2003.

[Omicini et al. 2004] A. Omicini, A. Ricci, M. Viroli, G. Rimassa: Integration Objective & Subjective Coordination in Multi-Agent Systems. *Proc. of the ACM Symposium on Applied Computing*, Nicosia, Cyprus, March 11–17, 2004.

[Phelps et al. 2003] J. Phelps, T. Wagner, V. Guralnik: TÆMS Agents for Distributed Emergency Response Coordination. Honeywell Labs TR ACS-R03-002, January, 2003.

[Raja & Lesser 2001] A. Raja, V. Lesser: Towards Bounded-Rationality in Multi-Agent Systems: A Reinforcement-Learning Based Approach. UMASS Dept. of CS TR-2001-034. 2001.

[Ricci et al. 2004] A. Ricci, M. Viroli, A. Omicini: Agent Coordination Contexts: From Theory to Practice. *in this volume*, 2004.

[Schumacher 2001] M. Schumacher: Objective Coordination in Multi-Agent System Engineering – Design and Implementation, LNAI 2039, Springer-Verlag, Heidelberg, 2001.

[Tambe & Zhang 2000] M. Tambe, W. Zhang: Towards flexible teamwork in persistent teams: extended report, in Special Issue: Best of ICMAS'98—Part II, *Autonomous Agents and Multi-Agent Systems*, **3**(2), 159–184, 2000.

[Viroli et al. 2004] M. Viroli, A. Ricci, A. Omicini: A Semantics for the interaction of agents with coordination artifact. *in this volume*, 2004.

[Wagner & Garvey 1997] T. Wagner, A. Garvey: Leveraging Uncertainty in Design-to-Criteria Scheduling. UMASS Department of Computer Science Technical Report TR-1997-011. 1997.

[Wagner et al. 1998a] T. Wagner, A. Garvey, V. Lesser: Criteria-Directed Task Scheduling. *Journal for Approximate Reasoning* – Special Scheduling Issue, **19**:91–118, 1998.

[Wagner et al. 1998b] T. Wagner, A. Raja, V. Lesser: Modeling Uncertainty and its Implications to Design-to-Criteria Scheduling, UMass Dept. of CS TR-1998-051. 1998.

[Wagner & Lesser 1999] T. Wagner, V. Lesser: Design-to-Criteria Scheduling: Real-Time Agent Control. UMASS Dept. of CS TR-1999-058. 1999.

[Wagner 2000] T. Wagner: Ph.D. Thesis: Toward Quantified Control For Organizationally Situated Agents. University of Massachusetts Amherst. February 2000.

[Wagner et al. 2001b] T. Wagner, J. Phelps, Y. Qian, E. Albert, G. Beane: A Modified Architecture for Constructing Real-Time Information Gathering Agents. in Wagner G., et al.(eds.): *Agent Oriented Information Systems (AOIS)*, pp. 121–135, 2001.

[Wagner et al. 2002] T. Wagner, V. Guralnik, J. Phelps: Software Agents: Enabling Dynamic Supply Chain Management for a Build to Order Product Line. Agents for Business Automation, 2002.

[Wagner 2003] T. Wagner: Tom's Scheduler Notes and Other Documentation. 16.04.2003. http://mas.cs.umass.edu/ wagner/research/.

[Wagner et al. 2003] T. Wagner, B. Horling, V. Lesser, J. Phelps, V. Guralnik: The Struggle for Reuse: Pros and Cons of Generalization in TÆMS and Its Impact on Technology Transition. in Satyadas A., Dascalu S.(eds.): *12th International Conf. on Intelligent and Adaptive Systems and Software Engineering (IASSE-2003)*, San Francisco, CA, USA, 2003.