PROGRESS REPORT
Reference period from 1 April 2004 to 30 September 2004

KUN, all

Revised version 31 May 2005

AMETIST DELIVERABLE 0.1.5

Project acronym: AMETIST
Project full title: Advanced Methods for Timed Systems
Project no.: IST-2001-35304
Project Co-ordinator: Frits Vaandrager
Project Start Date: 1 April 02
Duration: 36 months
Project home page: http://ametist.cs.utwente.nl/

Consortium

No	*Name*	*Short name*	*Country*
1	Katholieke Universiteit Nijmegen	KUN	NL
2	Robert Bosch GmbH	Bosch	D
3	Cybernetix Recherche	CYR	F
4	Axxom Software AG	Axxom	D
5	Terma A/S	Terma	DK
6	Aalborg University	AAU	DK
7	Universität Dortmund	Uni DO	D
8	VERIMAG	VERIMAG	F
9	Weizmann Institute of Science	WIS	IL
10	Laboratoire d'Informatique Fondamentale de Marseille	LIF	F
11	University of Twente	UT	NL

Contents

1 Industrial Objectives and Strategic Aspects

Ametist intends to contribute to solutions for the growing industrial need to design reliable and efficient time dependent systems. In particular, it intends to provide theory and tools for error-detection, control and optimisation of real-time distributed systems. Its approach will be based on translating state-of-the-art academic research into methods and tools that can be a basis for an industrial design practice of such systems.

In addition to its technological contributions, Ametist invests actively in knowledge transfer to the European industry of computer-aided timing analysis and design. Moreover, it is expected that the academic dissemination of the Ametist research results will influence and advance the field of timed systems research, and (indirectly) contribute to the education of future generations of system engineers.

Whereas timed automata and the tools for their analysis are widely accepted in academia and are being used at hundreds of universities and research laboratories all around the world, they have yet to find their way into industry. The aim of Ametist is to advance and mature the related models, tools, and methods to allow this situation to change.

The need for automatic tools that allow reasoning about time is evident. Beyond manufacturing, telecommunication and hardware, it is of essential importance for the growing market of embedded systems (from car electronics to home automation). However, there are several obstacles that seem to hinder the use of timed automata technology in industry at this time:

- Scalability: Currently, tools based on timed automata do not allow to handle big examples. There are industrial scale examples that have been treated with these tools but only after tedious manual simplification involving a lot of work in each case.
- Convenience: Current timed automata tools are stand-alone programs and their input formalisms lack important features for convenient specification in an industrial setting.
- Accessibility: To make optimal use of the currently available tools requires quite some sophistication on the user's part, which makes them practically inaccessible even to well-trained engineers.

Ametist aims at the (at least partial) elimination of these obstacles. The project moves towards this goal along several tracks. One is the treatment of real-life case studies from some candidate application domains to see if, indeed, the proposed models, tools and methodology are suited for them. Indeed much of the project's resources are being spent on case studies. A second direction aims to improve the situation regarding scalability, by introducing better algorithms and data-structures to model and manipulate large systems, in particular in the area of real-time controller synthesis, planning and scheduling. Moreover, the project aims at tool interaction to allow the interfacing of different tools, which can help to improve usability/convenience. The third track aims at synthesizing the accumulated results in order to assess the applicability of the project's vision and modify it according to feedback from the field.

2 Status, Achievements, Delays, Milestones

During the first month of the reporting period major effort was invested in producing the project deliverables for year 2, and the preparations for the second review meeting in Brussels on May 10. Following the review meeting, a Technical Evaluation Report was prepared by the experts Bo Wahlberg and Claude Le Pape (version 3.A, 08/06/04). This report and the recommendations contained in it were discussed extensively during the PCC meeting in Grenoble, France, on September 20, 2004 (see Section 3).

During the remaining part of the reporting period there were no specific milestones, and work has basically proceeded according to plan. Section 6 presents an overview of the scientific and technical results that were obtained by the project.

Major advances have been made in the area of tools. Several (new versions) of tools were released, implementing ideas that have been developed during the first two years of the project: UPPAAL, UPPAAL CORA, T-UPPAAL, S-Uppaal, TAOpt, IF-SCHED, DL-SAT, IF-DC, ELSE, and MOTOR. Tight connections and interfaces between all of these tools exist or are currently being developed. If we compare the current capabilities of these tools with what existed at the time AMETIST started, then it is fair to say that indeed we have moved the state-of-the-art to a new level of maturity (one of the main objectives of the project). Very recently, on October 23, Gerd Behrmann from AAU received the prestigious Spar Nord Foundation Prize (35.000 Euro) for his Ph.D. thesis [7] containing numerous contributions to the development of UPPAAL (see `www.ciss.dk` for more information (in Danish)).

Profiting from the new tools, major advances were also made with the AXXOM case study. Using UPPAAL CORA we succeeded to derive schedules that are competitive with those that were provided by industrial partner Axxom using its own tool Orion-Pi. Our experience with the AXXOM case study shows the application of model checking techniques for production scheduling is very promising. Still, considerable further work on modelling methods, reusebility of modelling patterns, identification and evaluation of heuristics, compasison with alternative approaches, all in the context of case studies of greater orders of magnitude, is needed to develop it into a readily applicable standard technique for scheduling.

3 Recommendations by the Experts

In this section, we discuss the recommendations from the Technical Evaluation Report based on the Second Review.

Recommendation 1

> *The work on stochastic approaches is very interesting, and should be further prioritized. It is important to implement this in the software tools.*

We agree that the stochastic analysis of the quality of schedules for the AXXOM case study carried out by partner UT [12] is very interesting. At UT currently much work is going on to further improve the MoDeST/MOTOR framework that has been used to carry out this analysis, and also some of the other academic partners (AAU, KUN, VERIMAG) are involved in research on stochastic models and analysis techniques. Nevertheless, we do not support the suggestion to further prioritize this work *within the context of* AMETIST. Tools for stochastic timed automata (such as MOTOR) have by far not reached the level of maturity of (regular/priced) timed automata tools such as UPPAAL, UPPAAL CORA and IF. Also there are great differences between the data structures used. Proper integration of say UPPAAL CORA and MOTOR is definitely beyond what can reasonably be expected from the AMETIST project. In fact, we believe that such an integation would be a very interesting topic for a successor project. We expect that for the remaining lifetime of AMETIST it is much more effective to focus on (1) consolidation and further integration of the new tools for (priced) timed automata that have been developed, (2) quantitatively assessing the progress made by continued application on challenging cae studies and benchmarks, and (3) reaching out to industry and the end-user community (cf Recommendations 6, 7 and 8). The latter activities will require most of our energy.

Recommendation 2

> *There is a need to develop better routines for writing deliverable reports, which contain more quantitative data and a structured analysis rather than just listing of results. In detail, the numerical results enabling to quantify progress were missing in D2.3b. The impact on usability, time scheduling and hybrid systems should be analyzed in greater detail.*

We will do our best. For clarification a number of comments can be made. Many productive researchers participate in the project, and numerous papers are produced on a variety of topics. Progress in timed automata technology is made on a broad frontier. Some of the progress can be quantified numerically, but much of the progress is of a qualitative nature (e.g., extension of the input language or GUI of a tool, the ability to model and analyse a certain application in a more "elegant" manner, etc). If we just look at the contributions made in a single year, they may seem unrelated and it is difficult to tell a coherent story involving all of them. The general picture only becomes clear if we consider the progress made over a number of years. Nevertheless, we think that there ought to be some place where all the results from a given reporting period are listed. Of coure, it is desirable to present quantitative data and a structured analysis whenever possible. Often, such information is contained in specific technical papers. In preparing the deliverables for year 2, we decided to give an overview of all tables with numerical results for the various tools in deliverable D2.5a. We did not want to repeat this information in D2.3.b, but a pointer to the tables in D2.5a would probably have helped to avoid confusion. We think it might be useful to not only provide the reviewers with progress reports (which by necessity contain listings of results) but also with a few selected technical papers that are representative for the progress made by the project.

Recommendation 3

> *A number of interesting case studies has been carried out which should be documented in a structured and complete way in D3.5.2 while emphasising the unique features developed.*

We propose that we try to give such a structured and complete overview in D3.5.3, referring to all additional case studies that have been carried out in the course of the project. If one looks at the miscelleneous case studies carried out over a longer period it becomes easier to give a structured overview of the areas in which timed automata technology has been applied, the progress that has been made, and the advantages and disadvantages of applying the technology in the different areas.

Recommendation 4

> *In D.1.2 give the reader some idea of the concrete possible applications of the new modelling methods in comparison with pre-existing methods and clearly describe links outside the project.*

We will provide an update of this deliverable.

Recommendation 5

> *The efforts in the case study of Bosch have been redefined from scratch. The new application scenario appears to be ambitious and rather complex for the remaining lifetime of the project. The consortium should decide within two months what specific tasks should be pursued and it should set the criteria to evaluate the progress.*

It was decided to devote a limited amount of resources from partners KUN and VERIMAG (in the order of 1PM) to assist partner Bosch in tackling this potentially very interesting case study.

Recommendation 6

> *A better tool strategy has to be developed. For the remaining period describe the realistic path and targets for tool synthesis and integration, including possible ways for the GUI implementation. It is very important to find ways to guarantee sustainable easy-to-use*

> *software tools for industry. The work on integration and convergence of tools needs to be improved.*

As should become clear from the paragraphs on Tasks 2.2, 2.4 and 2.5 in Section 6, much effort has been spend during the last 6 months to further improve and integrate the various tools from Ametist. In Section 7, we elaborate on the Ametist tool strategy.

Recommendation 7

> *Among the 4 major case studies identified in the project, only the Axxom study has received highest priority in the second year which provided sort of glue between the different WPs. However, more work is needed to be done on quantitative criteria for measuring the progress. After the initial experience from the other case studies decide at least on one more case study whose application is cross-cutting the various tools available in the consortium.*

The case studies proposed by Terma, Bosch and CYR, are all very interesting and instrumental to asses and further advance the state-of-the-art of timed automata technology. Hence we plan to continue to work on them. Nevertheless, we feel that these case studies are not suited to provide the same type of strong glue between the WPs as the AXXOM case study. Instead we suggest to consider a series of benchmark problems from the job shop literature, and to use these to obtain the quantitative data that will allow us to compare the different approaches developed within Ametist with each other and with tools that have been built by others. Partner Uni DO has proposed a list of more than 80 benchmarks from the job shop literature and made it available to the consortium. Also end-user panelist Carmen Consulting has proposed a number of interesting benchmarks to look at.

Recommendation 8

> *Maintain enough efforts to identify the end-user community and reach out to them by, for example, setting up an end-user panel and discussing application scenarios. Dissemination targeted at industry is crucial to the success of the overall goals.*

We agree. One of the reasons why we have asked for a three months extension of the project is that this will give us some more time – after completion of the research activities within Ametist — to reach out to industry and to disseminate the results of the project. Among the activities we are planning, a major one will be the organization of an Ametist conference for which all end-user panel members as well as other representatives from industry will be invited.

Recommendation 9

> *The consortium should explore ways and should define activities to use the results of the project in future strategic research projects like Networks of Excellence. A mid-term review is planned after m30 preparing the final period of the project and paving the way to exploit the results in related NoEs, i.e. HYCON, Artist2.*

The Ametist consortium has a strong presence within the Artist2 NoE, and there can be little doubt about the use of Ametist results within Artist2. Within Artist2, Kim G. Larsen is the leader of the cluster of Testing&Verification, see also [37]. Ed Brinksma is member of the Strategic Committee of Artist2.

Ametist partners UT and Uni DO participate in the HYCON NoE. Uni DO is particularly involved in HYCON workpackages on tool integration and development of methods for analysis and design of industrial hybrid systems. With respect to the latter, the search algorithm in TAOpt is related to the techniques that we intend to further develop for Optimal Control of Hybrid Systems within

HYCON. (The paper [49] is a shot in this direction). With respect to tools, the plan in HYCON is create a tool repository for hybrid systems. Within this activity, it may be reasonable to investigate in how far tools for timed systems (as produced in AMETIST) can or should be included.

Apart from Artist2 and HYCON, we would also like to build on the AMETIST in one or more smaller, more focused successor projects. Possible topics for such projects that we are currently considering include:

1. Full integration of tools/techniques for timed automata with tools/techniques for stochastic analysis (without loss of functionality/performance).
2. Use of timed automata technology for test generation.
3. Raising the level of timed automata tools even further, to enable their routine use in the practice of industrial design and value chain optimization. (this will require participation of an industrial tool builder)

We are open to suggestions for funding programs in which projects along these lines would fit.

Recommendation 10

> *To plan the exploitation routes, the consortium is strongly advised to establish and sign a consortium agreement that discusses and clarifies among the partners the Intellectual Property Rights (IPRs) issues.*

We will implement this suggestion.

4 Key Events During Reporting Period

During the reporting period two regular project meetings as well as the first review meeting took place:

- On May 10, 2004, the second review meeting took place in Brussels.
- The new case study proposed by Bosch was presented and discussed during a special project meeting in Nijmegen on June 9, 2004. This meeting was attended by 10 persons from Bosch, KUN and UT.
- On September 20-21, 2004, a project meeting was held in Grenoble, France, organised by VERIMAG. This meeting was attended by 25 persons.
- On september 22-24, 2004, the second FORMATS workshop took place in Grenoble, sponsored by AMETIST. Out of 69 submission, 24 papers were selected for presentation. There were 55 participants, including a substantial number of students. We had participants from Esterel Technologies, Honeywell, Sun Microsystems, SRI International and NASA.
- Partner UTorganized the highly successful first edition on the QEST (Quantitative Evaluation of Systems) international conference at the University of Twente in September 2004. The proceeding are published in the IEEE CS series [24].

The agendas and (most of) the slides for the above meetings are available on-line at `http://ametist.cs.utwente.nl/INTERNAL/MEETINGS/Meetings.htm`.

In addition to these meetings, several bilateral visits between partners took place, and AMETIST members attended many conferences where they presented the results of the project.

5 List of Deliverables

No	Description	Due Date	Delivery	Status	Resp Partner
4.4	Ametist Website	May 02	May 02	accepted	UT, all
0.1.1	Project Rep. - Progress & Evaluation	Oct 02	Nov 02	accepted	KUN, all
3.1.1	Case Study 1: Prel. Description	Oct 02	May 02	accepted	LIF, CYR
3.2.1	Case Study 2: Prel. Description	Oct 02	Apr 02	accepted	AAU, Terma
3.3.1	Case Study 3: Prel. Description	Oct 02	Sep 02	accepted	Uni DO, Bosch
3.4.1	Case Study 4: Prel. Description	Oct 02	Oct 02	accepted	Uni DO, Axxom
4	Dissemination and Use Plan	Oct 02	Oct 02	accepted (qual.)	VERIMAG, all
4.1.1	Ametist Workshop	Oct 02	Apr 02	accepted	VERIMAG
0.1.2	Project Rep. - Progress & Evaluation	Apr 03	Jun 03	accepted	KUN, all
0.2.1	Framework Report (v1)	Apr 03	Jun 03	accepted	VERIMAG, all
0.3.1	Financial Review	Apr 03	Jun 03	accepted	KUN, all
1.5	Modeling: Controller Synthesis	Apr 03	Apr 03	accepted	VERIMAG
2.3.a	A & T: State Space Representations	Apr 03	Jun 03	accepted	LIF
3.1.2	Case Study 1: Model	Apr 03	Jun 03	accepted	LIF, CYR
3.2.2	Case Study 2: Model	Apr 03	May 03	accepted	AAU, Terma
3.3.2	Case Study 3: Model	Apr 03	Jun 03	accepted	Uni DO, Bosch
3.4.2	Case Study 4: Model	Apr 03	Jun 03	accepted	Uni DO, Axxom
3.5.1	Misc. Case Studies: First Year Report	Apr 03	May 03	accepted	UT, all CRs
0.1.3	Project Rep. - Progress & Evaluation	Oct 03	Apr 04	accepted	KUN, all
0.1.4	Mid Term Assessment Report	Apr 04	May 04	accepted	KUN, all
0.2.2	Framework Report (v2)	Apr 04	May 04	accepted	VERIMAG, all
0.3.2	Financial Review	Apr 04	June 04	submitted	KUN, all
1.2	Modelling: Model Composition	Apr 04	May 04	accepted (qual.)	KUN
1.3	Modelling: Quantitative Modelling	Apr 04	May 04	accepted	UT
1.4	Modelling: Scheduling and Planning	Apr 04	May 04	accepted	Uni DO
2.1.1	A & T: Abstraction and Compositionality	Apr 04	May 04	accepted	KUN
2.2.1	A & T: Control Synthesis Algorithms	Apr 04	May 04	accepted	VERIMAG
2.3.b	A & T: State Space Representations (v2)	Apr 04	May 04	accepted (qual.)	LIF
2.4.a	A & T: Stochastic Analysis (v1)	Apr 04	May 04	accepted	UT
2.5.a	A & T: Tool Interaction (v1)	Apr 04	May 04	accepted	AAU
3.1.3	Case Study 1: Optimisation	Apr 04	May 04	accepted	LIF, CYR
3.2.3	Case Study 2: Optimisation	Apr 04	May 04	accepted	AAU, Terma
3.3.3	Case Study 3: Optimisation	Apr 04	May 04	accepted	Uni DO, Bosch
3.4.3	Case Study 4: Optimisation	Apr 04	May 04	accepted	Uni DO, Axxom
3.5.2	Misc. Case Studies: Second Year Report	Apr 04	Apr 04	accepted (qual.)	UT, all CRs
0.1.5	Project Rep. - Progress & Evaluation	Oct 04	Nov 04	submitted	KUN, all
0.1.6	Final Project Rep. - Progress & Evaluation	Apr 05	-	-	KUN, all
0.2.3	Framework Report (final)	Apr 05	-	-	VERIMAG, all
0.3.3	Financial Review	Apr 05	-	-	KUN, all
1.1	Modelling: Model Classification	Apr 05	-	-	VERIMAG
2.1.2	A & T: Structure Exploitation	Apr 05	-	-	KUN
2.2.2	A & T: Scheduling and Planning Algorithms	Apr 05	-	-	VERIMAG
2.3.c	A & T: State Space Representations (v3)	Apr 05	-	-	LIF
2.4.b	A & T: Stochastic Analysis (v2)	Apr 05	-	-	UT
2.5.b	A & T: Tool Interaction (v2)	Apr 05	-	-	AAU
3.1.4	Case Study 1: Final Report	Apr 05	-	-	LIF, CYR
3.2.4	Case Study 2: Final Report	Apr 05	-	-	AAU, Terma
3.3.4	Case Study 3: Final Report	Apr 05	-	-	Uni DO, Bosch
3.4.4	Case Study 4: Final Report	Apr 05	-	-	Uni DO, Axxom
3.5.3	Misc. Case Studies: Final Report	Apr 05	-	-	UT, all CRs
4.1.2	Ametist Conference	Apr 05	-	-	VERIMAG

6 Scientific and Technical Performance

Below we present a brief summary of the progress made within the different technical work packages/tasks by the different partners.

WP1: Modelling

In accordance with the original planning, the effort spend within this work package during the reporting period has been (relatively) small. Nevertheless, several relevant results were obtained within the context of AMETIST.

Task 1.1: Model Classification

VERIMAG formulated a general/unifying game-theoretic model for controller synthesis in which problems of scheduling under uncertainty can be formulated and solved [39].

Task 1.2: Model Composition

Partner KUN studied compositionality issues for probabilistic automaton models. Even though notions of (bi)simulation are known to be compositional for probabilistic systems, it turns out to be extremely difficult to obtain an abstract (trace based) compositional notion of behaviour for probabilistic systems. In a paper presented at ICTAC04 in China a partial solution was presented [19]. A switched probabilistic I/O automaton is a special kind of probabilistic I/O automaton (PIOA), enriched with an explicit mechanism to exchange control with its environment. Every closed system of switched automata satisfies the key property that, in any reachable state, at most one component automaton is active. We define a trace-based semantics for switched PIOAs and prove it is compositional. We also propose switch extensions of an arbitrary PIOA and use these extensions to define a new trace-based semantics for PIOAs.

For the final period of the project, KUN intends to work on (1) extension of the IOA toolset with real-time along the lines of [33, 34], (2) translation of (a subset of) the TIOA language into UPPAAL (which will help again to enlarge the user community for AMETIST), (3) study of assume/guarantee style verification rules, and (4) study of urgency/deadline predicates in the context of TIOA.

Task 1.3: Quantitative Modelling

An inventory of bisimulation and simulation relations, both weak and strong, for DTMCs and CTMCs was made, and their mutual relation investigated. Their logical characterizations in PCTL and CSL have been studied, as well as relations that exist between the models in the continuous and discrete time domains. This work is reported in [4].

For the remaining six months of the project, UT plans to complete the formal semantics of the MODeST modelling language, and to work on the issue of robustness of timed automata,

Tasks 1.4+1.5: Scheduling, Planning and Control Synthesis

The major activity of VERIMAG in the reporting period was the extension of the TA based methodology for scheduling to problems that involve conditional dependency between tasks, that is, the "result" of one task, known only after its termination, determines which tasks need to be further executed. This situations are common in program parallelization and in scheduling under uncertainty where the outcome of a task is not predetermined (e.g. possible failure). We have developed the model of the conditional precedence graph (CPG) and showed how it can be transformed, via TA, into a weighted game graph on which a problem of (worst-case) shortest path can be formulated and solved. We developed a forward search algorithm for this problem and managed to find solutions for problems with up to 100 tasks and 3-5 conditions. In this work we also moved away from breadth-first to depth-first search strategy. The results have been accepted to RTSS'04 [16].

WP2: Analysis and Tools

Task 2.1: Abstraction, Compositionality, and Structure Exploitation

S-UPPAAL A full version of the FORMATS'03 paper on symmetry reduction was prepared [26]. During the coming year we aim at integrating symmetry reduction as implemented in our prototype tool S-UPPAAL in the standard release of UPPAAL.

Task 2.2: Controller Synthesis and Scheduling Algorithms

Priced Timed Automata The model of priced timed automata (PTA) extends timed automata with a single cost variable, and allows a number of optimal scheduling problems to be formulated as cost-optimal reachability problems. The problem of cost-optimal reachability for PTA has been previously addressed within AMETIST [46] and constitutes the basis of the UPPAAL CORA branch of UPPAAL.

During the previous 6 months a number of new optimal scheduling problems in the context of priced timed automata has been considered.

Firstly, the substantially more difficult problem of synthesizing cost-optimal winning *strategies* in the presence of an adversary has been solved in [13]. The method described forms the basis of an algorithmic method for computing cost-optimal winning strategies. In [14] it is demonstrated how this method may be implemented in HyTech.

Infinite scheduling calls for strategies which enables a particular process to execute for ever. In this setting optimality may be formulated as strategy which minimizes the cost-time ratio in the long run. In [14] it is shown that the minimum limit cost-time ratio problem is in fact computable and that the optimal strategy (when one exists) is also computable.

More recently, we have considered conditional optimal reachability for multi-priced timed automata. More precisely, we consider the problem of determining the minimal cost of reaching a given target state, with respect to some *primary* cost variable, while respecting upper bound constraints on the remaining (secondary) cost variables. We show in [38] that this problem is computable using a zone-based algorithm.

In [10] we prove computability of the problem of parameter synthesis for the cost-bounded liveness problem for PTAs, that is the problem of synthesizing the *maximal* cost before a given target state is guaranteed to be reached. The problem has important applications to worst-case execution time analysis as demonstrated in paper on a number of task-graph scheduling problems where tasks have uncertain execution times.

Control Synthesis, a new generic algorithm Control synthesis aims to automatically compute a C controller for a plant P in order to achieve a control goal φ in the combined system: $P \times C \models \varphi$. Classical supervisory control theory looked at problems like avoiding undesirable states (safety). More recently, control synthesis was extended to the full range of branching time properties as presented by the μ-calculus. Partner LIF devised a new approach and generic algorithm for control synthesis of μ-calculus properties.

The dominant approach of the literature is to transform the problem $P \times ? \models \varphi$ into a specification of the controller, i.e. to construct from P and φ a specification $c(P, \varphi)$ and solve $? \models c(P, \varphi)$. Control synthesis is thus reduced to the satisfaction problem of then huge formulas.

The LIF approach instead transforms φ into a *controllability property* $P \models c(\varphi)$ which can be verified by a standard model checker. If the model checker yields that $P \models c(\varphi)$ holds, then we give a procedure for extracting from its computation a controller C.

The advantage of the approach is that $c(\varphi)$ has a low complexity for formulas written in a certain normal form and that the algorithmic complexity is entirely shifted to the well studied model checking problem.

An example application in AMETIST which will be carried out in the remaining time is the *faulty card problem* of the Cybernétix case study, see corresponding discussion.

Job-shop Scheduling In a previous deliverable, partner LIF showed how to use a partial order approach to improve the timed automata approach to (job-shop) scheduling problems. On the set of studied examples, the partial order approach reduced by a factor of up to 10 the number of states explored with the "best-first search" approach. However, the memory consumption of the best first search approach is a limiting factor to its applicability. Therefore, LIF studies how to combine the partial order approach with less memory hungry heuristic search methods. While there are encouraging results, this work will only be finished by the end of the year.

VERIMAG has improved its results on scheduling of tasks with uncertain but bounded durations. For a problem with 4 jobs, each with 6 tasks, and for randomly-drawn durations, our strategy was around 1.5% worse than a clairvoyant scheduling, compared to 4.5% for the conservative hole-filling strategy and 12.5% for the worst-case static schedule [1].

TAOpt A tool to realize the concept of combining graph search for timed automata with mixed-integer programming has been implemented by Uni DO. The tool named TAOpt currently allows to model arbitrary job shop scheduling problems by timed automata, and it solves the problems by embedding linear programming (LP) into the search of the reachability tree of the automata [44, 43]. The LP-step solves a relaxed (and automatically generated) MIP model, in order to obtain lower cost bounds that can be used to efficiently prune the search graph. In the past months, the focus was in particular on including existing and new techniques for reducing the search space, as well as on the development of efficient search heuristics:

- Techniques for partial order reduction, for weak and strong non-laziness, as well as for clock space reduction have been implemented.
- To improve the efficiency of solution, we have developed and implemented heuristics that reduce the state space by employing the notion of *similar* schedules (similar with respect to state sequences and clock evaluations).
- Various criteria to steer the search have been implemented (as, e.g., best lower bounds, cost-based priorities, and random components). Successful combinations could be identified in experiments.

The current version of TAOpt is ready to be used by end-users.

For the following six month, we plan to extend the modeling capabilities of TAOpt such that qualitatively different constraints and objective functions can be taken into account. In particular, we aim to implement modifications to consider restrictions on the resource-task assignments, change-over between resources, and parallel resource allocation. Extended cost functions will include costs for lateness, storage, and change-over.

Furthermore, we intend to provide interfaces to export the models used in TAOpt, thus to establish links to other tools (in accordance with Task 2.5). Algebraic models for the solution by MILP solvers can already be generated from the automaton specification in TAOpt. We investigate whether it is also feasible to create an XML interface for timed automata models, in order to connect to UPPAAL.

In a different effort, we have extended the concept of combining graph search for automata with embedded programming techniques to more complex systems than timed automata. As described in [48, 49], the approach can be used to synthesize (sub-)optimal controllers for hybrid automata with nonlinear continuous dynamics.

To assess the modeling strength and solution performance of TAOpt in comparison to alternative approaches, we have started numerical studies for different problem instances and benchmark problems. A set of over 80 scheduling benchmarks from literature has been proposed to the members of the consortium to be used for an extensive quantitative assessment. Initial experiments have shown that TAOpt can solve large scale problems (comprising 25 jobs, 15 tasks, and 15 resources) within 60 minutes with only a few percent of optimality gap.

Table 1 illustrates exemplarily for a set of smaller (synthetic) problems the solution performance of TAOpt [43]. The numbers reveal that considerably less nodes have to be explored when lower bounds (obtained from linear programming for a relaxed algebraic model) are employed, leading to much smaller memory requirements.

We intend to include a larger set of benchmark problems into the quantitative assessment, in order to evaluate for which parameterizations, search heuristics, and problem instances our tool is preferable over alternative approaches. A technical report will describe the results in detail.

Monitoring Real-Time Properties VERIMAG started to investigate monitoring as a lightweight alternative to verification of large systems. The idea is to specify a desired property in some dialect of real-time temporal logic, and generate from it a property monitor that observes the behavior of the system in question (either a simulation or the real system) and alerts the user when the property is violated. We have developed and implemented a backward monitoring procedure for a bounded fragment of the logic MITL and are currently working on forward procedure and on extensions to other specification formalisms such as timed regular expressions. [40].

Real-Time Testing A number of contributions of applying the timed automata technology to the planning of real-time testing has been made. In [36] methods for on-line real-time conformance testing with respect to a given timed automaton model is provided and implemented in the tool T-UPPAAL (branch of UPPAAL). Also, optimal off-line testing of real-time systems with guaranteed coverage of the specification, is obtained using optimal reachability of timed automata [29, 30].

At VERIMAG, new techniques for conformance testing of real-time systems have been developed and applied to the K9 Mars Rover case-study provided by NASA [35, 11].

Table 1: Number of nodes explored with different search schemes: BFS - Standard best first search; BLB - Best lower bound first search using embedded linear programming; for BLB 8, BLB 5, and BLB 3, embedded linear programming is only applied to every 8-th, 5-th, or 3-rd layer of the search tree.

Operations	Jobs	BFS	BLB 8	BLB 5	BLB 3	BLB
2	2	18	19	16	13	11
2	3	92	87	72	47	41
2	4	467	270	264	243	141
2	5	3017	3009	2783	2593	2136
2	6	12789	10146	7501	5364	3984
3	2	40	31	23	18	16
3	3	275	102	48	33	33
3	4	3037	790	597	437	293
3	5	26020	12491	8862	7140	5154
3	6	227208	92692	71903	56220	39496

Task 2.3: State Space Representations

Partial Order Semantics The previously reported work by partner LIF on a partial order semantics approach to the reachability problem of timed automata has been further consolidated: a long version of the article with detailed proofs was produced and Sarah Zennou has finished her thesis report on this work. The defense of the thesis is planned for December.

In addition to this, we currently conduct experiments with several variants of the zone list approach for the representation of symbolic states of timed automata. The goal is to have a faster test for zone inclusion due to more compact data structures. This work is going on, results will be reported at the end of year three.

DL-SAT VERIMAG continues to explore SAT solving to difference logic as an alternative computational engine to timed verification and optimization. We have developed a new SAT solver, DL-SAT, which integrates "learning" and other procedures to improve performance. On certain classes of problems, including job-shop scheduling, the results were indeed better than in other solvers developed recently [20].

Task 2.4: Stochastic Techniques

MOTOR At UT continued work was done on the development of the tool MOTOR for integrated functional and non-functional analysis of embedded systems. An application case study was carried out analysing a distributed device-absence checking protocol using the stochastic simulation facilities of MOTOR. This work is reported in [32, 28]. The tool has also been applied for the analysis of the AXXOM case study (see below). The work on MOTOR also belongs to Task 2.5.

Work was done by UT on the symbolic modelling of probabilistic systems through the use of rational power series. This approach has the advantage that in a certain cases exact solutions can be obtained that do not suffer from accumulated errors of numerical approximations. This work is reported in [21].

ETMCC During the final period of the project, UT intends to work on algorithms for time-bounded reachability for CTMDPs, and on improving the performance of the ETMCC Markov chain model checker by about one order of magnitude.

Task 2.5: Tool Interaction

UPPAAL UPPAAL version 3.4.7 has been released on October 5, 2004, and is now fully compatible with Java JRE 1.5 (beta 1).

A new UPPAAL tutorial note fully compatible with version 3.4.7 has been written as a contribution to a Summerschool on Formal Methods for Real-Time Systems (Bertinoro September 2004) [9].

Before the end of the year UPPAAL version 3.5.0 will be made available allowing the user to define complex functions using a fully integrated imperative programming language (syntactically this language is very similar to C). We expect this new feature will make it substantially simpler to make natural models of real-time systems containing complicated discrete data-structures (e.g. routing tables, ready queues, etc.).

UPPAAL CORA Currently, there is a complete version of UPPAAL CORA implementing the entire modelling language of priced timed automata[1]. That is, timed automata models can be augmented with a single real-valued cost variable with different growth rates in locations and discrete cost increases on transitions. The type of reachability problems solvable by UPPAAL CORA is finding the minimum cost of reaching some goal location.

[1] UPPAAL CORA may be downloaded from `http://www.cs.aau.dk/~behrmann/cora/`.

The search can be guided with respect to two internal meta variables called *heur* and *remaining* that are associated with each state and which, respectively, states an ordering among states to be explored and an underestimate on the remaining cost of reaching the goal from a given state. The search strategies supported by UPPAAL CORA are depth-first, random depth-first, best depth-first, breadth-first, smallest/largest *heur* first, and best-first. Pruning with respect to *remaining* is applied to the first three strategies.

Furthermore, UPPAAL CORA has been extended with U-code, allowing the user to define function in a C-like language and use these when modelling automata. Usually, components such as queues are modelled as automata, but now, U-code enables the user to specify the queue as an internal array that can be modified by automata calling appropriate functions.

UPPAAL CORA has been applied to the AXXOM lacquer production case study by the Twente and Nijmegen partners.

Currently, in collaboration with the industrial partner Carmen Systems, a comparison study is being conducted on the competitiveness of the (priced) timed automata approach to scheduling on a traditional operations research benchmark. More specifically, the benchmark of vehicle routing with time windows which can be stated as: Given a depot of goods, a fleet of vehicles with limited capacity, and a set of dispersed customers with associated demands and time windows, find a schedule that assigns to each vehicle a sequence of customers such that the total demand of the customers does no exceed the capacity of the vehicle and starting from the depot, the vehicle can visit the customers in turn within their respective time windows and return to the depot. Solutions to the vehicle routing problem will be evaluated with respect to various optimality criteria such as the total travel distance and the finishing time of each vehicle.

T-UPPAAL On May 16, 2004, the version 1.1 of the testing tool T-UPPAAL was released[2] and presented at FATES'04 [36].

T-UPPAAL is a testing tool for black-box conformance testing of embedded and real-time systems. Given a formal timed automata model of the system under test (SUT) and its assumed operating environment, T-UPPAAL automatically generates and executes timed test sequences. T-UPPAAL is an on-line testing tool which means that it continuously executes test events on the SUT as they are being generated, and events from the SUT is checked against the model. The observed behavior is required to be timed trace included in the specification.

The tool is based on the UPPAAL engine which is a model-checker of real-time systems modeled as networks of timed automata.

The system under test is attached to T-UPPAAL via a test-adapter (an SUT specific software layer) and considered as a black-box since its states cannot be directly observed; only communication events via input/output channels. The user supplies T-UPPAAL with the closed timed automata network of SUT model in parallel composition together with assumptions on environment.

The explicit environment model is an important feature: it can be used to generate test events only that are realistic it the operating environment. It may also be fully-permissible, meaning that the environment (testing tool in this case) can offer any input at any moment and accept any output at any moment. Finally it can be used to guide the test (which is randomized) to produce particularly interesting behaviors.

In addition off-line methods for real-time testing using the optimal reachability features of UPPAAL (CORA) has been developed and implemented [29, 30].

IF VERIMAG worked on the IF tool set which can now be used for verifying certain dialects of UML (part of the work done within the Omega project) [42, 15].

ELSE The prototypical timed automata tool ELSE developed at LIF has received a new intermediate representation. This representation is aimed to be sufficiently expressive for easy translation

[2] See http://www.cs.auc.dk/ marius/tuppaal/

from both IF and UPPAAL and their recent extensions (like U-code). A front end reading IF and representing it in the ELSE intermediate representation is nearly complete. A front end for UPPAAL is planned. Currently, the existing ELSE tool chain is being rewritten for the new intermediate representation.

This intermediate format is also designed for easier transformation of of specifications in intermediate representation into code for analysis with zones (like UPPAAL and the previous version of ELSE), but also to generate formulas for bounded model checking.

The architecture of ELSE consists of three parts, transformation into intermediate representation, transformations on the intermediate representation, and code generation. When finished, this will allow a list of tool chains made of combinations of tools in the AMETIST project, some examples:

- UPPAAL or IF source code, analyzed with ELSE back-end (to be finished within AMETIST).
- UPPAAL or IF source code + reachability property, translated to *difference logic* and analyzed with VERIMAG's difference logic solver (to be finished within AMETIST).
- UPPAAL or IF source code + reachability property, translated to boolean satisfaction problem and analyzed with state of the art SAT solvers (probably not within AMETIST).
- UPPAAL or IF source code, transformed within ELSE intermediate representation, analyzed with UPPAAL back-end (no resources attributed yet).

WP3: Case Studies

Task 3.1: Cybernétix Case Study

In the last six months, LIF and CYR have been working on the faulty card problem of the HPX machine. As a reminder, the problem is that faulty cards may be detected only during programming and must be replaced. The replacement is difficult because the conveyor belt is essentially filled with cards on their trajectory through the system. Cybernétix has devised a scheme for rejecting and replacing faulty cards which engineers believe to be robust against an arbitrary number of faults. This is important since faulty cards are due to production problems and typically arrive in bursts.

Two goals were pursued, the *verification* of the Cybernétix approach, and a solution by *control synthesis*.

For the verification of the Cybernétix approach we have remodeled the plant and the controller in IF. Compared to previous work, we have found an *improved abstraction* that significantly reduces the number of states that the model has to go through. The abstraction is so strong that in the error free case of the Super Single mode the card identities have no influence on the state space, i.e. the model then reduces to "material flow" (this becomes wrong in the case with errors).

The contribution is an optimized IF model of both the machine and the controller with the Cybernétix faulty card mode. This work is close to being finished.

Subsequently, we intend to explore our new control synthesis algorithm to the faulty card problem. For the desired property "the machine produces a correctly ordered flow of cards" we should be able to compute a control algorithm for the same order of magnitude as for the previously executed synthesis of the control algorithm for the error free case.

Also at WIS some work has been done on the Cybernétix case study, which has been reported in [25].

Task 3.2: Terma Case Study

The work on the Terma case study is continuing along two independent directions. In [23] the bit-state hashing technique of UPPAAL has been applied to automatically obtain the (buffer-size) scheduling principle originally proposed (and manually verified) by G. Weiss [53].

An International Master Thesis student is currently working on completing and verifying an extended model of the memory interface with more details wrt. the organization of the used SDRAM into account.

Task 3.3: Bosch Case Study

During the reporting period, Boschworked out the proposal for a new case study. The new case study is concerned with modelling component behaviour of an airbag ECU and proving correctness and safety properties. Bosch has worked out a set of component models in UPPAAL and a set of correctness and safety requirements for system behaviour.

The proposed case study was presented and discussed during a special project meeting in Nijmegen on June 9, 2004, and during the AMETIST Meeting in Grenoble on September 20-21. It was decided to devote a limited amount of resources from partners KUN and VERIMAG (in the order of 1PM) to assist partner Bosch in tackling this potentially very interesting case study.

During the remainder of the project, Boschwill further investigate the usability of timed-automata modelling and verification in safety-critical embedded system development. Bosch also intends to investigate approaches for automatic test set generation from UPPAAL component models.

Task 3.4: AXXOM Case Study

Considerable effort has been invested by several partners in the further elaboration of the AXXOM case study.

The initial case study was extended by Axxom by adding more restrictions and more products and orders to the model. This extension was made available in April 2004 and can be viewed on the Ametist home page. Since then Axxom has supported the academic partners working on the case study by answering questions and explaining this case study. In order to do stochastic calculations on the availability factor, Axxom provided a complete new case study based on a packaging facility. This case study can also be viewed on the Ametist home page.

The first successful analysis using timed automata and UPPAAL ignored two parameters of the original problem statement, viz,. the so-called performance and availability parameters, which relate to the occurrence of failures, repairs, and other unforseeable events. Both the UPPAAL model was extended to include these features, and analysed, also considering larger numbers of orders. In a collaborative effort of partners UT, KUN and AAU, also costs for storage and delays have been taken into account, which required a switch to cost-optimal scheduling using UPPAAL CORA. Using UPPAAL CORA we succeeded to derive schedules that are competitive with those that were provided by industrial partner Axxom using its own tool Orion-Pi. These results are reported in [8]. For the remaining period of the project, we plan to obtain realistic schedules for the AXXOM case study including the (problematic) restriction on working hours, and increase the number of jobs in the case study by one or two orders of magnitude.

At UT, also a stochastic analysis was carried out using the MoDeST/MOTOR framework. Using the discrete event simulation engine of the integrated Moebius tool the quality of the schedules was assessed with respect to timeliness, utilization of resources, and different kinds of reliability assumptions. This was done using a stochastic timed automaton (derived from the original UPPAAL model) to reflect the occurrence of unplanned events such as machine breakdown etc. This work is reported in [12].

The algebraic model, developed by Uni DO in year 2 for solving the AXXOM case study by MIP, was extended in order to account for the modifications in the newest version of the lacquer production scheduling problem. The efforts have addressed the following points:

- Additional constraints to model change-over costs and change-over times, as well as the interruption of production during night and on weekends have been introduced.
- The preemption of tasks has been included in the algebraic formulation.

- The object function has been modified to include costs for materials and the allocation of resources.
- The case study has been solved by MIP within 40 minutes of computation time for an instance with 14 jobs, where the preemption of tasks and the interruption of production at night was excluded [45].

So far, no solution could be obtained for problem formulation with preemption of tasks and breaks at night. Modifying the algebraic model to efficiently solve the problem for this case remains an issue for the upcoming weeks. In addition, we aim to find a (near-)optimal solution for the problem instance with 29 jobs. In addition to use MIP for the solution, the tool TAOpt (see subsection on Task 2.2) will be applied to this case study.

Also VERIMAG is working on the new version of the AXXOM case study, but there are no results yet.

For the final deliverable, we intend to compare the different approaches taken within the Ametist project (UPPAAL CORA, TAOpt, MILP, ORION-Pi and posibly VERIMAG) by applying them to exactly the same version of the AXXOM case study, and to several of the the job-shop benchmarks.

Task 3.5: Miscellaneous Case Studies

ASML case study For a case-study of a wafer scanner from the semiconductor industry, partner KUN showed how model checking techniques can be used to compute (i) a simple yet optimal deadlock avoidance policy, and (ii) an infinite schedule that optimizes throughput. Deadlock avoidance is studied based on a simple finite state model using SMV, and for throughput analysis a more detailed timed automaton model has been constructed and analyzed using the UPPAAL tool. The SMV and UPPAAL models are formally related through the notion of a stuttering bisimulation. The results were obtained within two weeks, which confirms once more that model checking techniques may help to improve the design process of realistic, industrial systems. Methodologically, the case study is interesting since two models (and in fact also two model checkers) were used to obtain results that could not have been obtained using only a single model (tool). The results were presented at the 1st International Symposium on Leveraging Applications of Formal Methods (ISoLA2004) [27] and are referred to in patent application ASML ref. P-1784.010.

Biphase mark protocol The biphase mark protocol is frequently used for communication at the physical level of the ISO/OSI hierarchy. An important property of the protocol is that bit strings of arbitrary length can be transmitted reliably, despite differences in the clock rates of sender and receiver, and variations of the clock rates (jitter), and distortion of the signal after generation of an edge. In [52], partner KUN shows how the protocol can be modelled naturally in terms of timed automata. UPPAAL is used to derive the maximal tolerances on the clock rates, for different instances of the protocol, and to support the general parametric verification that was formalized using the proof assistant PVS. Based on the derived parameter constraints instances of BMP are proposed that are correct (at least in the model) but have a faster bit rate than the instances that are commonly implemented in hardware.

7 Tool Strategy

Figure1 presents an overview of the tools that have been used/developed within Ametist. We have classified these tools along two dimensions: (1) the total number of Ametist PM's spent on their development during the whole lifetime of the project, and (2) the "maturity" of the tool on a scale of 0 (academic prototype, minimal user interface, for internal use only), to 1 (industrial, commercial, good support, widely used). All numbers are estimates. Assessing maturity, of course, is a difficult exercise. As becomes clear from the figure, there is a spectrum ranging from tools such as UPPAAL, which have a very nice GUI, are easy to use, and although academic have almost

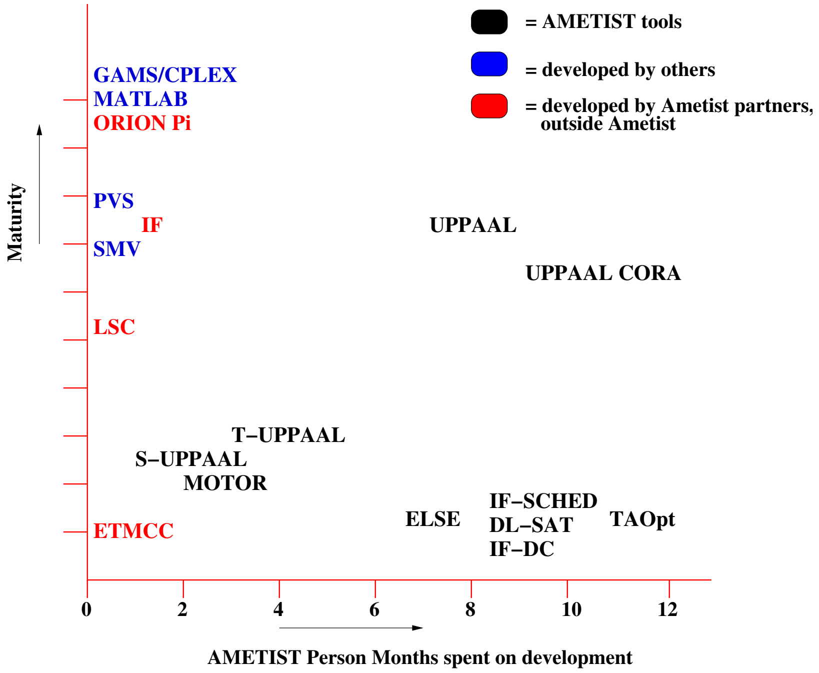

Figure 1: Overview of the tools developed/used by AMETIST.

industrial quality, to highly experimental tools such as ELSE which are primarily intended to test the viability of some new verification approach.

Below we comment on our vision for the different tools. These comments are complementary to the comments made about these tools in the previous section.

- **S-UPPAAL** The tool S-UPPAAL will cease to exist when the notion of symmetry reduction that it supports has been fully integrated within UPPAAL.
- **T-UPPAAL** In terms of maturity, T-UPPAAL is ranked much lower than UPPAAL. Time testing is a hot topic nowadays and tools such as T-UPPAAL have much potential. Further development of T-UPPAAL could be a good topic for a successor project. Only a small fraction of the PMs actually spent on the development of T-UPPAAL has been supported by AMETIST.
- An integration of TAOpt and CPLEX exists in the following respects: (a) job shop scheduling problems modelled in TAOpt can be automatically generated in the input language of CPLEX, and solved with the latter by MILP; (b) to generate lower cost bounds in the TAOpt algorithm, CPLEX is called repeatedly (on-the-fly).
- **IF** Before the start of AMETIST, VERIMAG has developed the IF tool-set for verification of complex large system models which include timed automata as a subset. It employs both discrete time and dense time analysis methods. The tool is very mature, has front-end connections with input languages such as SDL and UML, objects, model reduction methods, etc. Most of the developments in this tools are being done outside the AMETIST project. The activities in the AMETIST project are more exploratory in nature and are realized by prototype independent tools (here referred to as IF-SCHED, DL-SAT and IF-DC) that may profit from the know-how and the libraries of IF but are not integrated into it unless they become more mature.
- **IF-SCHED** Early experiments with job-shop scheduling using IF convinced us that the “classical” zone-based algorithms constitute an overkill for many scheduling problems. Consequently VERIMAG has developed an independent point-based tool for finding shortest path

in timed automata and used it to solve the job-shop problem, the task-graph problem and the preemtive job-shop problem. Note that for preemptive scheduling we use stopwatch automata that typically complicate zone-based algorithms, but cause no problem for our point-based method. In all the abovementioned problems our tool's performance matched that of state-of-the-art tools described in the academic literature.

Later we have extended the tool in two directions. The first was the treatment of temporal uncertainty in task duration. The approach that we took, dynamic programming for timed automata, required the use of zones and used the DBM library of IF. Since IF is limited to forward analysis it could not be used by itself for this problem. The second direction is the introduction of discrete uncertainty (conditional precedence graph) and the resolution of the associated problem using search algorithms on game graphs.

All these tools are rather rudimentary in terms of user interface as most of the effort is in improving the algorithmics. The scheduling problem is specified in a classical way (job-shop, task graph, conditional precedence graph) and transformed automatically into timed automata. Some of the results of this work has been transferred to IF: it is now possible to add cost and use search algorithms for finding optimal runs or runs that satisfy certain cost constraints. Non-lazy schedules which are the basis of our point-based methods were also implemented in IF, although not in a direct way. The first results on the Axxom case-study have been obtained using this extension of IF.

Our strategy for the future is to build a specialized tool for scheduling applications based on all the abovementioned algorithmic results, with a richer schedule description language that can cover more complex real-life constraints such as the one manifested in the new Axxom case study. We hope to have some results within the project lifetime but most of the effort is planned for future projects.

- **DL-SAT** We have invested a lot of effort in trying to find alternative and fully-symbolic methods for timing related problems in order to tackle the state explosion problem. Inspired by the relative success of SAT-based methods for purely discrete systems, VERIMAG developed a series of solvers for difference logic, a Boolean logic with timing constraints. It is hoped that the particular character of timing constraints will contribute to more efficient algorithms compared to classical methods that treat mixed discrete and arbitrary linear constraints.

  Our first solver, MX-Solver served for us as an introduction to the domain and through which we have explored many dead ends. Last year we started the development of a new solver, DL-SAT, into which many of the lessons of the first solver have been incorporated. Several performance improvement have been obtained, but so far no breakthrough has been achieved. The solver beats some other recent solvers developed elsewhere, especially in scheduling problems, but at the moment it is inferior compared to the classical timed automata methods that we use in IF and in the schedule-specific tools, and of course, compared to commerical tools such as CPLEX. We intend to keep on working on it, and if the long awaited breakthrough arrives, this will be a great news for timed automaton verification.

- **IF-DC** An alternative method for coping with state explosion is divide-and-conquer, that is to decompose the system into components and create a conservative approximation of each components having a smaller number of states and clocks than the concrete systems. We developed an automatic abstraction method for acyclic closed systems and now we are working on its extension to open systems. This work currently focuses on timed automata derived from timed digital circuits (because of their relative simplicity and regular structure) but it could be extended naturally to more general types of timed automata. The tool has a front-end that translate digital circuits expressed in common industrial format (SDF) into IF timed automata, and uses the exploration and minimizations methods of the latter to generate the abstractions.

  If successful, the techniques behind IF-DC can probably be integrated in tools such as IF and UPPAAL, but this will only happen after AMETIST has ended.

- **LSC** We did not integrate the LSC tool with any of the timed automata based tools, but our modeling examples together with advances in TA technology show the applicability of such a tool. We identified areas (e.g. production line planning) in which an integrated tool would be desirable.

  The translation of LSCs to timed automata is straightforward but there are many semantical issues that should be resolved. A good the way to proceed is to define the semantics that is best suited for applications. In this respect the work at WIS is very relevant because it gives examples that may direct future developments.

  At AAU, currently two MsC students are working on a translation of LSCs to (UPPAAL style) timed automata.

- **MOTOR** and **ETMCC**. Only a very small part of the work on these tools has been done on these tools within Ametist. Experience with the AXXOM case study shows that integration of these tools with the other Ametist tools is potentially very useful. Carrying out the integration and raising the maturity level of these tools is an excellent topic for a successor project.

- **ELSE**. The goal of ELSE is to explore the potential of a new verification approach for timed automata based on partial order semantics. If the outcomes of experiments with ELSE (to be carried out during the lifetime of Ametist) are sufficiently promissing, then on the long run full integration of the ELSE approach in the UPPAAL or IF tools would be a sensible option. Within Ametist we will develop translations from UPPAAL and IF to ELSE, in order to facilitate benchmarking.

- **UPPAAL (CORA)** UPPAAL CORA is becoming nearly as mature as UPPAAL, and within short time they will share most of the internal code-base and GUI. Several more man months have been invested in developing UPPAAL and (to a lesser extent) UPPAAL CORA than those declared within Ametist due to the collaboration of AAU with Uppsala. From all the tools that have been (further) developed within Ametist, UPPAAL (CORA) is currently the only "sustainable easy-to-use software tool for industry" (cf Recommendation 6 by the experts). During the final months of the project, we intend to work hard to disseminate our results, especially to industry.

- **UPPAAL** versus **IF**. On the long run integration (close coupling) of UPPAAL and IF appears to make a lot of sense. However, one should realize that the two tools evolved by and large independently, and their philosophy is rather different. Whereas UPPAAL has been designed as a stand alone tool, the IF language and toolset has been developed as an intermediate format that allows one to link high level specifications in languages such as UML and SDL to formal verification tools.

  Linking the two tools is difficult as the IF language is considerably more expressive. In addition to abstract data types, IF can also include directly C/C++ code in specifications. Moreover, in IF transitions can have arbitrary sequential code (including loops and conditionals). Only by the end of this year a new version of UPPAAL will be made available that fully integrates a C-like language. The time model of IF based on urgency types (instead of invariants) is also richer (more compositional). IF also has dynamic creation of processes and there are priorities. Finally, a difference is that in IF communication is by asynchronous message passing and processes have a message queue, whereas UPPAAL uses synchronous communication. Although none of the differences is insurmountable, actual coupling of UPPAAL and IF will be a time consuming activity. For Ametist, the benefits of having the translation were not that big, so we decided not to work on this.

We intend that by the end of the project the graph with tools as vertices and edges whenever tools can (partially) talk to each other, is connected. Further integration of all our tools within the lifetime of Ametist is neither realistic nor desirable. What we *do* plan to do is to make an extensive comparison of the tools that aim at solving scheduling problems (UPPAAL CORA, IF-SCHED, ELSE, TAOpt, ORION Pi, GAMS/CPLEX) by considering versions of the AXXOM case study and benchmarks taken from the literature on job-shop scheduling. Attacking these case studies with different tools and approaches is a step towards integration because it shows the strengths end weaknesses of every tool and where they can complement each other.

8 Management, Co-ordination, Resources

There are no major deviations from the workplan for any of the partners or any of the workpackages during the reporting period, apart from those that were discussed already in Del 0.1.4. In this section, we will discuss the planning for the remaining 9 months of the project and in particular the consequences for the workplan of (a) the reallocation of money from partner Terma to other partners, (b) the allocation of money from the central coordination budget to support additional person months for 5 partners, (c) the extension of the project with 3 months.

8.1 Reallocation of Money

As indicated already in Deliverable 0.1.4, we expect that all partners will spend their budget as specified in the contract, with exception of Terma. Tomas Hune, the AMETIST representative at at Terma, got a new job as section leader with focus on new and uniform development methodologies to be used throughout Terma. Since this cannot be combined with work on AMETIST and no other employee within Terma is able to take over, the PCC has agreed to reallocate personnel costs. The new budget for personnel costs for Terma will be EURO 9691 (instead of the original EURO 57.691). Since Terma operates under the FC model half of the remainder, i.e. EURO 24.000, will become available for reallocation between the other partners.

In the AMETIST central coordination budget EURO 80.000 has been reserved for "Other costs". This money has not been touched during the first two years of the project and the PCC agreed to make EURO 70.000 of it available for reallocation. The total travel and subsistence budget for partner KUN is EURO 142.050. Out of this money, EURO 24.000 is for KUN as partner and the remaining EURO 118.050 is intended to cover general project costs related to project meetings, travel outside western Europe, etc. The current estimate is that this money will be necessary and sufficient for these purposes.

The PCC agreed to use the total amount of 94.000 EURO that is available for reallocation to increase the budget for personnel costs of five partners as follows:

Participant	*Total of Estimated Eligible Costs*	*Maximum Community Contribution*	*PM*
KUN	22125	22125	4
AAU	22125	22125	4
Uni DO	5500	5500	1
VERIMAG	44250	22125	6
UT	44250	22125	4
Total	138250	94000	19

The partners will use the additional money to get the following work done:

- Uni DO will invest in comparative studies of the tools and algorithms developed and publication of results.
- A significant amount of time (larger than foreseen at the start of the project) will be needed by partner KUN to coordinate the preparation of the final deliverables and the review. In addition, the completion and documentation of the Axxom case studies will require a larger investment than planned originally.

- AAU intends to work on:
  1. Improvement and evaluation of UPPAAL CORA (integration of Ghant Charts, implementation of LU-abstraction, implementation of a distributed version).
  2. Improvements and evaluation of (classic) UPPAAL (slicing and abstractions wrt discrete (C-code) part of the model, integration of Live Sequence Charts, experimental evaluation of UPPAAL on NORDUGRID as part of the work on Test & Verification in the ARTIST2 Noe).
- VERIMAG intends to use the additional money to work on the following topics:
  1. Extending non-lazy scheduling to problems with relative deadlines.
  2. Exploring more combinations of depth-first and best-first search in job-shop scheduling.
  3. DL-SAT and BMC for timed automata: a new version of DL-SAT including a new translation to CNF, new and more efficient encoding of BMC, and more.
  4. Extending conditional scheduling with probabilities.
- UT wants to get out of the Axxom case study everything possible (3 persons are working on this). Also, UT will continue to have a relative large share in dissemination activities (in particular related to case studies).

8.2 Resource Planning

All figures in the table below denote persons months (PMs). All PMs figures correspond to the additional PMs as mentioned in the CPFs; the own contributions of partners that work under the AC model are not included. Ideally, for each partner the sum of the resources used in years 1+2 and the planned resources for the remainder of the project (m25-39) should equal what has been planned, i.e., the sum of the number of PMs in the Annex and the number of PMs that has been reallocated. In reality, the two sums are not always equal, as becomes apparent from the last column ("Deviation"), which lists for each partner the difference between the two sums.

	Used in Yr 1+2	*Planned m25-36*	*Planned m37-39*	*Result Reallocation*	*Total Planned (Annex 1)*	*Deviation*
KUN	37.6	23.0	8.4	4.0	65.0	0
Bosch	1.9	0.9	0.2	0	3.0	0
CYR	15.5	8.2	2.0	0	28.9	- 3.2
Axxom	7.5	3.5	0.3	0	11.3	0
Terma	0.9	0.2	0.1	-5.5	6.7	0
AAU	20.2	26.0	1.0	4.0	42.4	0.8
Uni DO	24.0	12.0	3.0	1.0	36.0	2.0
VERIMAG	85.0	10.0	5.0	6.0	78.0	16.0
WIS	27.7	9.4	0.6	0.0	30.6	7.1
LIF	45.7	12.0	2.0	0	63.0	-3.3
UT	40.4	26.0	3.0	4.0	39.9	25.5
Total	306.4	131.2	25.6	13.5	404.8	44.9

The (significant) deviations in the last column can be explained as follows:

- For CYR, the total number of PMs turned out somewhat lower than planned: the hourly rates were higher than planned due to the participation of project and product managers in Ametist. For similar reasons, the total number of PMs at WIS turned out somewhat higher than planned.

- Partners VERIMAG and UT are investing much more in the project than promissed. In fact, VERIMAG and UT invested more PMs in AMETIST in the first 2 years than planned originally for 3 years.
- The PhD student working on AMETIST in LIF left the university in April, after successfully defending her thesis, and no replacement was found for the last few months.

Note that for partners KUN, VERIMAG, and Uni DO the figures for "Used in Yr 1+2" are slightly different from the preliminary figures mentioned in Deliverable 0.1.4 [2].

The table below exhibits planned resources per partner per work package for the final period of the project, i.e., m25 - m39.

	WP0	*WP1*	*WP2*	*WP3*	*WP4*	*Planned in Period*	*Total Used Yr 1+2*	*Total Planned (Annex 1)*
KUN	6.0	6.0	3.0	15.0	1.4	31.4	37.6	65.0
Bosch	0.1	0	0	1.0	0	1.1	1.9	3.0
CYR	0.1	0	6.5	3.5	0.1	10.2	15.5	28.9
Axxom	0.1	0	0	3.6	0.1	3.8	7.5	11.3
Terma	0.1	0	0	0.2	0	0.3	0.9	6.7
AAU	1.0	2.0	16.0	6.0	2.0	27.0	20.2	42.4
Uni DO	1.0	2.0	6.0	5.0	1.0	15.0	24.0	36.0
VERIMAG	2.0	3.0	6.0	1.0	3.0	15.0	85.0	78.0
WIS	0.1	3.0	3.7	3.0	0.2	10.0	27.7	30.6
LIF	0.7	1.0	10.0	2.0	0.3	14.0	45.7	63.0
UT	3.0	4.0	5.0	15.0	2.0	29.0	40.4	39.9
Planned in Period	14.2	21.0	56.2	55.3	10.1	156.8		
Total Used Yr 1+2	19.8	62.9	132.2	85.5	6.0		306.4	
Total Planned (Annex 1)	33.0	72.3	175.9	107.6	16.0			404.8

In our work we intend to follow the recommendations made by the reviewers:

- Our plan is to finalise the technical work by m36 while devoting the last 3 months of the project, i.e. months 37-39, solely to documentation, dissemination and exploitation, in particular to prepare the final AMETIST workshop in June 2005.
- Most of the work will focus around the case studies from Axxom and Cybernetix; the investment of AMETIST resources in other case studies will be limited.
- We will only invest AMETIST resources in those tools for which consolidation and validation is feasible within the lifetime of the project.

References

[1] Y. Abdeddaïm, E. Asarin, and O. Maler, *Scheduling with timed automata*, Theoretical Computer Science (to appear), 2004.

[2] AMETIST, *Periodic progress and management report for period from 1 april 2003 to 31 march 2004*, May 2004, Deliverable 0.1.4 from the IST project AMETIST.

[3] C. Baier, B. Haverkort, H. Hermanns, J.-P. Katoen, and M. Siegle (eds.), *Validation of stochastic systems: A guide to current research*, LNCS, vol. 2925, Springer-Verlag, 2004, 467 p. Tutorial volume.

[4] C. Baier, J.-P. Katoen, H. Hermanns, and Verena Wolf, *Comparative branching-time semantics for markov chains*, Tech. Report CTIT-TR-04-32, University of Twente, 2004, Submitted for publication. 64 pages.

[5] N. Bauer, S. Engell, R. Huuck, S. Lohmann, B. Lukoschus, M. Remelhe, and O. Stursberg, *Verification of plc programs given as sequential function charts*, Integration of Software Specification Techniques for Applications in Engineering, LNCS, vol. 3147, Springer, 2004, pp. 517–540.

[6] N. Bauer, R. Huuck, B. Lukoschus, and S. Engell, *A unifying semantics for sequential function charts*, Integration of Software Specification Techniques for Applications in Engineering, LNCS, vol. 3147, Springer, 2004, pp. 400–418.

[7] G. Behrmann, *Data-structure analysis for formal verification*, Ph.D. thesis, Aalborg University, 2003.

[8] G. Behrmann, E. Brinksma, M. Hendriks, and A. Mader, *Scheduling lacquer production by reachability analysis - a case study*, Extended abstract submitted to the IFAC World Congress 2005. Full version submitted to Workshop on Parallel and Distributed Real-Time Systems 2005.

[9] Gerd Behrmann, Alexandre David, and Kim G. Larsen, *A tutorial on uppaal*, Formal Methods for the Design of Real-Time Systems, Lecture Notes in Computer Science, no. 3185, Springer Verlag, 2004, pp. 200–236.

[10] Gerd Behrmann, Kim G. Larsen, and Jacob I. Rasmussen, *Beyond liveness: Efficient parameter synthesis for time bounded liveness*, To appear, 2004.

[11] S. Bensalem, M. Bozga, M. Krichen, and S. Tripakis, *Testing conformance of real-time applications by automatic generation of observers*, Runtime Verification (RV'04), 2004.

[12] H.C. Bohnenkamp, H. Hermanns, R. Klaren, A. Mader, and Y.S. Usenko, *Synthesis and stochastic assessment of schedules for lacquer production*, Proc. QEST'04, LNCS, sep 2004, To appear.

[13] Patricia Bouyer, Franck Cassez, Emmanuel Fleury, and Kim G. Larsen, *Optimal strategies in priced timed game automata*, Proceedings of the 24th Conference on Fundations of Software Technology and Theoretical Computer Science (FSTTCS'04) (Chennai, India), Lecture Notes in Computer Science, Springer-Verlag, 2004, To appear.

[14] ______, *Synthesis of optimal strategies using HyTech*, Proceedings of the Workshop on Games in Design and Verification (GDV'04) (Boston, Massachusetts, USA), Electronic Notes in Theoretical Computer Science, Elsevier Science Publishers, 2004, To appear.

[15] M. Bozga, S. Graf, L. Mounier, and I. Ober, *IF tutorial*, SPIN'04 Workshop on Model-Checking of Software, Barcelona, Spain, April 2004.

[16] M. Bozga, A. Kerbaa, and O. Maler, *Optimal scheduling of acyclic branching programs on parallel machines*, 2004, Submitted for publication.

[17] E. Brinksma, *Testing times: On model-driven test generation for non-deterministic real-time systems*, In Proceedings Fourth International Conference on Application of Concurrency to System Design (ACSD 2004), IEEE CS Press, 2004, pp. 3–4.

[18] E. Brinksma, T. Krilavivcius, and Y. Usenko, *A process algebraic approach to hybrid systems*, Submitted to the IFAC World Congress 2005. 6 pages.

[19] L. Cheung, N.A. Lynch, R. Segala, and F.W. Vaandrager, *Switched probabilistic I/O automata*, Proceedings First International Colloquium on Theoretical Aspects of Computing (ICTAC2004), m Guiyang, China, 20-24 September 2004, 2004, To appear. Full version available as Technical Report NIII-R0437, NIII, Radboud University Nijmegen, September 2004.

[20] S. Cotton, E. Asarin, O. Maler, and P. Niebert, *Some progress in satisfiabilty checking for difference logic*, Proc. of FORMATS-FTRTFT 2004, 2004, Submitted for publication.

[21] C. Daws, *Symbolic and parametric model checking of discrete time Markov chains*, Proc. 1st International Colloquium on Theoretical Aspects of Computing (ICTAC'04) (Guiyang, China), September 2004, To appear in LNCS.

[22] S. Engell, A. Maerkert, G. Sand, and R. Schultz, *Aggregated scheduling of a multiproduct batch plant by two-stage stochastic integer programming*, Optimization and Engineering **5** (2004), 335–359.

[23] Juhan Ernits, *Abstraction based analysis and arbiter synthesis: Radar memory interface card case study revised*, Proceedings of Nordic Workshop on Programming Theory 2004, 2004.

[24] G. Franceschinis, B.R. Haverkort, J.-P. Katoen, and M. Woodside (eds.), *Quantitive evaluation of systems proceedings*, IEEE CS Press, 2004.

[25] David Harel, Hillel Kugler, and Gera Weiss, *Some methodological observations resulting from experience using lscs and the play-in/play-out approach*, Tech. Report MCS04-06, The Weizmann Institute of Science, 2004, Submitted to the post-conference proceedings of the Dagstuhl Seminar 03371 on Scenarios: Models, Algorithms and Tools.

[26] M. Hendriks, G. Behrmann, K.G. Larsen, P. Niebert, and F.W. Vaandrager, *Adding symmetry reduction to Uppaal*, Proceedings First International Workshop on Formal Modeling and Analysis of Timed Systems (FORMATS 2003), *September 6-7 2003, Marseille, France*, LNCS, vol. 2791, Springer Verlag, 2004, Full version available as Technical Report NIII-R0407, NIII, University of Nijmegen, February 2004.

[27] M. Hendriks, N.J.M. van den Nieuwelaar, and F.W. Vaandrager, *Model checker aided design of a controller for a wafer scanner*, Report NIII-R0430, Nijmegen Institute for Computing and Information Sciences, University of Nijmegen, 2004, To appear in em Proceedings ISOLA'04.

[28] H. Hermanns, H.C. Bohnenkamp, D.N. Jansen, J.-P. Katoen, and Y.S. Usenko, *An industrial-strength formal method: A modest survey*, In 1st International Conference on International Symposium on Leveraging Applications of Formal Methods, (ISOLA) (Paphos, Cyprus), LNCS, Springer, 2004, Invited contribution.

[29] Anders Hessel, Kim G. Larsen, Brian Nielsen, Paul Pettersson, and Arne Skou, *Time-Optimal Test Cases for Real-Time Systems*, 3rd International Workshop on FORMAL APPROACHES TO TESTING OF SOFTWARE (FATES 2003) (Montral, Qubec, Canada), October 2003, In affiliation with the 18th IEEE International Conference on AUTOMATED SOFTWARE ENGINEERING (ASE 2003).

[30] ———, *Time-Optimal Test Cases for Real-Time Systems–extended abstract*, 1st International Workshop on Formal Modeling and Analysis of Timed Systems (FORMATS), 2003, Invited Talk by Paul Pettersson.

[31] D.N. Jansen and H. Hermanns, *Dependability checking with stocharts: Is train radio reliable enough for trains?*, In 1st International Conference on Quantitative Evaluation of Systems (QEST), IEEE CS Press, 2004, pp. 250–259.

[32] J.-P. Katoen, Henrik Bohnenkamp, H. Hermanns, and J. Klaren, *Embedded software analysis with motor*, LNCS, pp. 268–294, Springer, Bertinoro, Italy, 2004.

[33] D.K. Kaynar, N.A. Lynch, R. Segala, and F.W. Vaandrager, *A framework for modelling timed systems with restricted hybrid automata*, Proceedings of the 24th International IEEE Real-Time Systems Symposium (RTSS03), m December 3-5, 2003, Cancun, Mexico, ieeepress, 2003, pp. 166–177.

[34] ______, *The theory of timed I/O automata*, Tech. Report MIT-LCS-TR-917, MIT Laboratory for Computer Science, Cambridge, MA, 2003.

[35] M. Krichen and S. Tripakis, *Black-box conformance testing for real-time systems*, 11th International SPIN Workshop on Model Checking of Software (SPIN'04), LNCS, vol. 2989, Springer, 2004.

[36] Kim Larsen, Marius Mikucionis, and Brian Nielsen, *Online Testing of Real-time Systems using Uppaal*, International workshop on Formal Approaches to Testing of Software (Co-located with IEEE Conference on Automates Software Engineering 2004, Linz, Austria.) (Jens Grabowski and Brian Nielsen, eds.), September 2004.

[37] Kim G. Larsen and Brian Nielsen, *ROAD MAP on Hard Real-Time Development Environments. Chapter 7: Tools for Verification and Validation*, Year 1 deliverables of Project IST-2001-34820 ARTIST:Advanced Real-Time Systems ARTIST IST-2001-34820, Aalborg University, May 6 2003, Road map is to be publised as a book/volume by Springer Verlag.

[38] Kim G. Larsen and Jacob I. Rasmussen, *Optimal conditional reachability for multi-priced timed automata*, To appear, 2004.

[39] O. Maler, *On optimal and sub-optimal control in the presence of adversaries*, In Proc. WODES'04, 2004, Invited talk at WODES'04.

[40] O. Maler and D. Nickovic, *Monitoring temporal properties of continuous signals*, In Proc. FORMATS-FTRTFT 2004, 2004.

[41] M. Massink, J.-P. Katoen, and D. Latella, *Model checking dependabiliy attributes of wireless group communication*, Dependable Systems and Networks - Performance and Dependability Symposium (DSN 2004) (Firenze, Italy), IEEE CS Press, 2004, pp. 711–720.

[42] Iulian Ober, Susanne Graf, and Ileana Ober, *Model checking of UML models via a mapping to communicating extended timed automata*, SPIN'04 Workshop on Model Checking of Software, 2004, vol. LNCS 2989, 2004.

[43] S. Panek, O. Stursberg, and S. Engell, *Job-shop scheduling by combining reachability analysis with linear programming*, Proc. 7th Int. Workshop on Discrete Event Systems, 2004, pp. 199–204.

[44] ______, *Optimization of timed automata models using mixed-integer programming*, Formal Modeling And Analysis of Timed Systems, LNCS, vol. 2791, Springer, 2004, pp. 73–87.

[45] Sebastian Panek and Sebastian Engell, *Scheduling of a pipeless multi-product batch plant using mixed-integer programming combined with heuristics*, Proc. ESCAPE 15, ESCAPE 15, 2004, PO-151.

[46] J. Rasmussen, K. G. Larsen, and K. Subramani, *Resource-optimal scheduling using priced timed automata*, Proc. 10th Int. Conf. of Tools and Algorithms for the Construction and Analysis of Systems (TACAS'2004), Lecture Notes in Computer Science, vol. 2988, Springer-Verlag, 2004, pp. 220–235.

[47] M.P. Remelhe, S. Lohmann, O. Stursberg, S. Engell, and N. Bauer, *Algorithmic verification of logic controllers given as sequential function charts*, Proc. IEEE Conf. on Computer-Aided Control System Design, 2004.

[48] O. Stursberg, *Dynamic optimization of processing systems with mixed degrees of freedom*, Proc. 7th Int. Symposium on Dynamics and Control of Process Systems, 2004, p. ID: 164.

[49] ———, *A graph-search algorithm for optimal control of hybrid systems*, Proc. 43rd IEEE Conf. on Decision and Control, 2004, to appear.

[50] O. Stursberg, A. Fehnker, Z. Han, and B.H. Krogh, *Verification of a cruise control system using counterexample-guided search*, Control Engineering Practice **12** (2004), no. 10, 1291–1303.

[51] J. Till, S. Engell, S. Panek, and O. Stursberg, *Applied hybrid system optimization - an empirical investigation of complexity*, Control Engineering Practice **12** (2004), no. 10, 1269–1278.

[52] F.W. Vaandrager and A.L. de Groot, *Analysis of a biphase mark protocol with Uppaal and PVS*, Tech. Report MIII-R0445, NIII, Radboud University Nijmegen, 2004.

[53] Gera Weiss, *Optimal Scheduler for a Memory Card*, Research report, Weizmann, 2002.