Merging Planning, Scheduling & Verification - A Preliminary Analysis

Amedeo Cesta(1), Alberto Finzi(2), Simone Fratini(1), Andrea Orlandini(3), Enrico Tronci(4)

(1) *ISTC-CNR*
Via S.Martino della Battaglia 44, I-00185 Rome, Italy
Email: name.surname@istc.cnr.it

(2) *DSF "Federico II" University*
Via Cinthia, I-80126 Naples, Italy
Email: finzi@na.infn.it

(3) *DIA "Roma TRE" University*
Via della Vasca Navale 79, I-00146 Rome, Italy
Email: orlandin@dia.uniroma3.it

(4) *DI "La Sapienza" University*
Via Salaria 198, I-00198 Rome, Italy
Email: tronci@di.uniroma1.it

Designing Artificial intelligence (AI) planning and scheduling systems suitable for supporting human mission planners in their daily work is an increasing and challenging research stream at ESA.

One of the key points to take into account for raising an effective introduction of AI based systems is the acceptability of these technologies for the end users. In fact, automated planning and scheduling systems often brings solutions to the users which are not "obvious" and immediately acceptable for them. This is due to the fact that these tools are able to take into account quite an amount of temporal and causal constraints and the employed resolution processes are often designed to optimize the solution with respect to non trivial evaluation functions. Thus, the study of tools for verifying and validating plans and schedules produced by AI systems might help in facilitating the introduction of these technologies.

This paper presents a preliminary analysis of the issues concerned with the application of NuSMV and UPPAAL, two software tools for the formal verification of finite state systems, to the validation of the solutions produced by MrSPOCK, the "Mars Express Science Plan Opportunities Coordination Kit", a recent effort for building a timeline based planning tool within the ESA project APSI (Advanced Planning and Scheduling Initiative). MrSPOCK addresses the Mars Express spacecraft long-term plan optimization problem, a challenging multi-objective optimization problem which goal is to build pre-optimized skeleton plan that allocates spacecraft activities for science, maintenance and communication for uplink/downlink The preliminary analysis in this paper aims at presenting the approach we have followed as well as discussing the problems encountered and the lessons learned in trying the automatic validation and verification of MrSPOCK's solutions against the temporal and causal constraints in the symbolic model that drives the tool in building the plans.

INTRODUCTION

The use of Autonomous Planning Systems (APS) presents novel and difficult testing challenges that traditional software does not face. In fact, APS are able to automatically generate a wide range of different plans, i.e. proposing different possible solutions, and automatic manipulation of system level constraints is facilated. Thus, APS must be shown to generate correct plans and the planning models used still need to be verified. In this sense, validation of planning models and solutions has been studied in several works[1]. For instance, in [14] a verification and validation approach is used and discussed within the Remote Agent Experiment context. In [16, 12] Livingstone and HSTS domain models are validated exploiting model checking techniques. In [19], formal verification is used in order to check the existence of undesirable plans with respect to the domain model. While, VAL [10] is a plan validation tool for PDDL that was successfully used during the International Planning Competition since 2002.

---

[1] See also a specific workshop at ICAPS-05.

Current AI planning literature shows how timeline-based planning can be an effective competitor for classical planning to tackle complex domains which require the use of both temporal reasoning and scheduling features (see [15, 11, 6, 18]). The work described here is connected to timeline planning because of a general effort to build a reusable software framework for modeling space missions problems using timelines (see [3]). The timeline-based approach models the P&S problem by identifying a set of relevant *features* of the planning domain which need to be controlled to obtain a desired temporal behavior. Timelines model entities whose properties may vary in time and which represent one or more physical (or logical) subsystems which are relevant to a given planning context. The planner/scheduler plays the role of the controller for these entities, and reasons in terms of *constraints* that bound their internal evolutions and the desired properties of the generated behaviors (goals).

In our current work, we plan to explore different perspectives in the integration of V&V with timeline based planning and scheduling techniques. The long term goal is to obtain a software environment in which both technologies are integrated and the application developers may take advantage of the co-existence of the two tools while knowledge engineering new application. In fact, we are investigating the use of model checking to assist in the validation of models used in such a systems and the verification of solvers exploited.

Moreover, one of the key points to take into account to foster effective introduction of AI P&S systems in real world is to develop end user trust in the related technologies. Automated P&S systems often brings solutions to the users which are neither “obvious” nor immediately acceptable for them. This is due to the ability of these tools to take into account quite an amount of temporal and causal constraints and to employ resolution processes often designed to optimize the solution with respect to non trivial evaluation functions. To increase technology trust, the study of tools for verifying and validating plans and schedules produced by AI systems might be instrumental.

This paper presents a preliminary report of the issues concerned with the use of two software tools for formal verification of finite state systems to the validation of the solutions produced by MrSPOCK, a recent effort for building a timeline based planning tool in an ESA project. In particular, we consider (a) the application called MrSPOCK (Mars Express Science Plan Opportunities Coordination Kit) as planning system; and (b) two prominent software tools in formal verification, NuSMV and UPPAAL.

As a first step towards the desired general environment, in the present work we focus our attention on the validation of the solutions produced by MrSPOCK. Such a task provides a tool to check whether the used models reflects correctly the needs of end users. Notwithstanding the initial extent of the current effort we have already received some interesting hints on the whole activity.

The preliminary analysis in this paper aims at presenting the approach we have followed as well as discussing the problems encountered and the lessons learned in trying the automatic validation and verification of MrSPOCK’s solutions against the temporal and causal constraints in the symbolic model that drives the tool in building the plans.

TARGET SCENARIO

As said in the introduction we have focused our current activity in developing a validation tool for the MrSPOCK planner recently developed within the APSI project for MARS EXPRESS long, medium and short term planning. The open problem was to support the collaborative problem solving process between the science team and the operation team of the space mission. These two groups of human planners iteratively refine a plan containing all activities for the mission. The process starts at the long term plan (LTP) level – three months of planning horizon – and is gradually refined to obtain fully instantiated activities at short term plan (STP) level - one week of planning horizon. This process continuously leads to weekly STPs, which are then further refined every two days to produce final executable plans. Goal of MrSPOCK has been to develop a pre-planning optimization tool for spacecraft operations planning. Specifically, we have focused on the generation of a *pre-optimized skeleton LTP* which will then be subject to cooperative science team and operation team refinement (see [2] for a detailed description of the addressed problem).

We observe that, at the first step of the negotiation process, one source of approximation comes from the fact that the operation team has only partial information about the requested science operations for MARS EXPRESS. Payload Operation Requests (PORs), that is the requested science operations. Reference to such requests is given only when the Pointing Timeline Request is issued by the science team on the basis of the input skeleton plan generated by operation team. Indeed, these science requests often require the satellite to point to the planet, reducing its ability to obtain energy from the Sun and to send data back to Earth. Also, science operations consume power and exclude the possibility of performing maintenance operations. On the other hand, the operation team requires the spacecraft to perform maintenance and other service maneuvers in order to maintain the spacecraft operational. Overall, the two groups of managers have to cooperatively converge on a plan which resolves the complex interplay between science, pointing direction, power, data transmission and maintenance operations.

In this context, the challenge of MrSPOCK is to provide an automated procedure for producing a *good* skeleton plan, i.e., a LTP that takes into account the needs of both parties, thus reducing the effort in reaching an agreement on a medium-term plan – one month planning horizon. Overall, the generated LTP should be such that: (a) the number of (expensive) iterations between science and operation team is reduced; (b) a set of objective functions – the total volume of data for down-link operations; the number of pericentres for science operations; the number and the uniform distribution of uplink windows – are optimized.

Problem Required Constraints

For each orbit followed by the spacecraft, the baseline operations are split into three phases: (1) around the *pericentre* (the orbital closest to the target planet); (2) around the *apocentre* (the orbital more far away from the planet); (3) *between* the pericentre and apocentre passages. Around pericentre, the spacecraft is generally pointing to the centre of the planet thus allowing observations of the planet surface – generically referred to as *Science operation*. Between pericentre and apocentre passages, the spacecraft is generally Earth pointing for transmitting data to Earth. *Communication* with Earth should occur within a *ground station availability window*. Ground station visibility can either partially overlap or fully contain a pericentre passage. Additionally, *Maintenance* operations should occur around the apocentre passages.

At present, given these requirements, an initial skeleton plan for MARS EXPRESS is generated by the operation team by allocating over the planning horizon (which generally covers hundreds of orbits) three different types of decisions: selection of the *Maintenance* windows (generally centered around the apocentre events and used primarily for *momentum wheel-offloading*); selection of the *Communication* windows among the set of available ground stations (the so-called ground stations de-overlapping); selection of the windows for *Science* operations around pericentre events.

Additionally there are quite an amount of *hard* and *soft* constraints to be satisfied. Constraints on uplink windows frequency and duration require four hours uplink time for each 24 hours (soft constraint). Moreover, there should be given the possibility to split a four-hour uplink window in two two-hour uplink windows. Apocentre slots for spacecraft maintenance windows must be allocated between 2 and 5 orbits apart, and the maintenance duration is of 90 minutes to be centered around the apocentre interval.

Communication activities are source of several temporal constraints to be considered as hard. For example: (1) the minimum/maximal durations for the X-band transmitter in the *on* state, (2) the minimum duration for the X-band transmitter in the state *off*; (3) the periods in which the X-band transmitter has to be *off* (e.g., eclipses, occultations, slewing manoeuvres and non-Earth pointing status).

Furthermore, there are preferences to follow for ground station selection (called *de-overlapping* in mission terminology). Ground stations have different features like different dish diameters (there are 70 meters dish antennas, 35 meters and 34 meters). Usually, they allow both uplink and downlink communications, but there are cases where it is only possible to downlink. Additionally, there are ground stations owned by different agencies and they should be used according to some policy restriction.

A Hybrid Solver

MrSPOCK uses the modeling capabilities of a quite general software framework, named TRF (Timeline-based Representation Framework), which provides the basic elements for modeling the relevant entities in the space contexts (see [2] for an introduction to the TRF's architecture and for a MrSPOCK's more detailed description). The TRF is the synthesis of our long experience as both researchers in P&S and developers of working tools for ESA. The opportunity to investigate in this direction has been provided by the APSI (Advanced Planning and Scheduling Initiative) project, an ESA initiative to develop an open framework for the flexible support of mission planning systems.

In building MrSPOCK, we have followed a hybrid approach. We have used (1) the timeline representation and management features of the TRF for the problem representation, (2) a general purpose planning and scheduling system, called Open Multi-Component Planner and Scheduler (OMPS [7]), built on top of the TRF, for planning the timelines and (3) a domain dependent specific solver that guarantees the satisfaction of the problem's constraints not modeled in the domain description and that performs genetic driven plans optimization, exploiting the domain independent underlying planning system.

The TRF is designed as a layered architecture: there is an underlying temporal database (that provides primitives to represent and manage time points and temporal constraints), a timeline management and representation layer above the temporal database (that provides primitives to represent temporal flexible plans as timelines) and an upper level that provides a unified, shared representation of the plan. The OMPS planner and scheduler provides a portfolio of domain independent planning and scheduling procedures on top of the TRF. MrSPOCK uses the TRF and OMPS and proposes a domain dependent search procedure on top of the TRF.

In MrSPOCK, we use the TRF features by representing the domain with two different types of timelines (see Figure 1): (1) *Controllable State Variables*, which define the search space of the problem, and whose timelines ultimately represent the solution to the problem; (2) *Uncontrollable State Variables*, representing values imposed over time which can be only observed. Here, we use a single controllable state variable to model the spacecraft's pointing mode, which specifies the temporal occurrence of science and maintenance operations as well as the spacecraft's ability to communicate. The values that can be taken by this state variable, their durations (represented as a pair $[min, max]$) and the allowed transitions among them are synthesized by the automaton in the right in Figure 1. This is an effective way to capture part of the set-up necessary constraints in the solution.

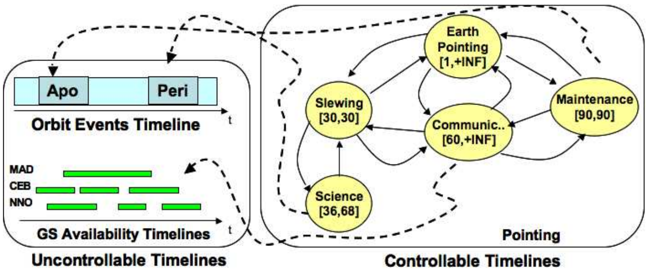

Figure 1: The domain model

In addition, we instantiate two uncontrollable state variable to represent contingent events such as orbit events and communication opportunity windows. One state variable type component maintains the temporal occurrences of pericentres and apocentres ("PERI" and "APO" values on the timeline in Figure 1, left, top) of the spacecraft's orbit (they are fixed in time according to the information found in an orbit events file), while the other state variables maintains the visibility of three ground stations ("MAD","CEB" and "NNO" timelines in Figure 1, left, bottom). These state variables have as allowed values {Available(?rate,?ul_dl,?antennas), Unavailable()}, where the ?rate parameter indicates the bitrate at which communication can occur, ?ul_dl indicates whether the station is available for upload, download or both, and the ?antennas parameter indicates which dish is available for transmission. Any valid plan needs temporal synchronizations among the pointing timeline and the uncontrollable timelines (represented as dotted arrows in Figure 1): science operations must occur during Pericentres, maintenence operations must occur during Apocentres and communications must occur during ground station visibility windows. In addition to those synchronization constraints, the pointing mode timeline must respect the transition among values specified by the automaton and the minimal and maximal duration specified for each value (in the automaton as well).

On top of this model, MrSPOCK's solver allocates science, maintenance and communication activities and exploits the TRF features for synchronizing them with orbit events and ground station visibilities. By querying the TRF that maintains the actual temporal occurrence of previously allocated activities, orbit events and available visibility windows, the domain dependent planner allocates new activities enforcing constraints on uplink windows and maintenance frequencies. The choice between science, maintenance and communication activities (when available stations allow a communication activity) is driven by a chromosome. Once the complete temporal plan has been instantiated, it is measured and used in a genetic optimization loop aiming at maximizing science allocation.

MODEL CHECKING

Aiming at fostering users to act more confidently with MrSPOCK and introducing an effective supporting tool, we perform an additional effort to verify solutions validity with a different technology. In fact, it is worth noting that performing an independent model verification can provide support also during design phase, helping in validating the domain model. Furthermore, when a hybrid approach that mixes different solving procedures is followed, as in MrSPOCK, the need of an indipendent solution verification process increases.

In fact, since not all the problem constraints are taken into account in the domain description of the general planning system, the proof of correctness of the domain independent planning system can not ensure the correctness of the produced solutions. Moreover a genetic optimization process is performed within MrSPOCK and the chromosome management performed by the genetic algorithm might invalidate the solution as well. Hence the need of verifying the solutions of such a kind of hybrid solver.

In this sense, V&V techniques represent a needed complementary technology in developing domain independent architectures for automated problem solving. In particular, model checking is a well known technology used to verify requirements and design for a variety of real-time embedded and safety-critical systems. Moreover, it is particularly well suited for exploring the relevant execution paths of systems with multiple processes running in parallel. Then, two prominent software tools in model checking, NuSMV and UPPAAL, are considered, focusing on the problem of the validating solutions produced by MrSPOCK.

NuSMV [4] is a model checker for temporal logics. It has a dedicated modelling language, which permits the definition of concurrent finite state systems in an expressive, compact, and modular way. The SMV specification uses variables with finite types, grouped into a hierarchy of module declarations. Each module states its local variables, their initial value and how they change from one state to the next. The properties are expressed in Computation Tree Logic (CTL). CTL is a branching-time temporal logic, which means that it supports reasoning over both the breadth and the depth of the tree of possible executions.

UPPAAL [13], whose acronym comes from joining the names of UPPsala and AALborg universities that built it, is a tool box for modeling, simulation, and verification of real-time systems. The verifier covers the exhaustive dynamic behavior of the system for proving safety and bounded liveness properties. A UPPAAL model consists of a set of timed automata, a set of clocks, global variables and synchronizing channels. A node in an automaton may be associated with an invariant, for enforcing transitions out of the node. An arc may be associated with guards, for controlling when this transition can be taken. On any transition, local clocks may get reset and global variables may get re-assigned. Channels are used in order to synchronize transitions on different automata. Analogously to NuSMV, verified properties are stated in CTL.

PLAN VALIDATION WITHIN MrSPOCK

In this section, a possible mapping of planning models into model checkers models is presented and discussed. In particular, in Figure 2, an integrated validation architecture is presented. This mapping allows us to perform different validation tasks: domain model validation can be performed by just considering the timelines model description and verifying safety or liveness properties similarly to [12]; furthermore, once we extend the model checker model with plan descriptions, plan validation can also be deployed to verify the correctness of solutions proposed. In the following, we focus our attention on the latter validation type.

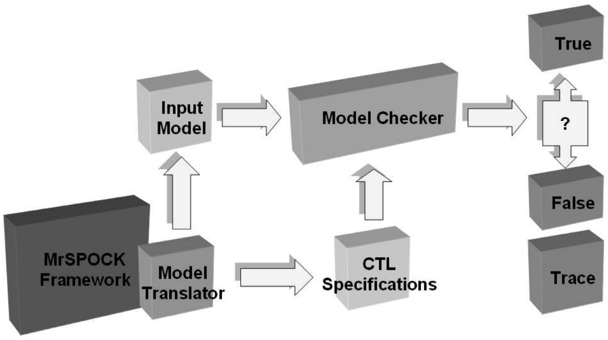

Figure 2: Plan validation architecture exploiting model checking.

In order to perform plan validation, MrSPOCK domain and plan are encoded in a new model, while a property assuring plan validity is defined. Then, both the model and the property are provided as inputs to model checking tool. In this sense, an initial goal is to build the new model, such that it can be encoded both in NuSMV and UPPAAL input languages, and state a suitable property for plan validity check.

Starting to address such an issue, target application domain description and a completely specified temporal solution plan, generated by MrSPOCK, are to be encoded in the new model. Such a model is to be translated in the model checkers input languages. Thus, plan validity property is defined according to such a model.

In order to simplify the presentation, let us consider in the following no parameters handling. Although such a feature is managed in our modelling, it is not quite relevant with respect to the current discussion.

Modelling MrSPOCK Domain and Plan

First, domain and plan are to be encoded in the new model. Within the framework, a domain model is represented as a set of components (i.e. timelines). For each of them, a set of *consistency features* are defined in order to state the nominal component's behaviors. Consistency features represent transition and temporal constraints related to one singular timeline. Looking at Figure 1, consistency features related to the spacecraft's pointing mode are, for instance, the duration constraint on Science activity (within $\{36, 68\}$ interval) or the transitions from *Maintenance* to *Earth* or *Comm* states.

In addition, some inter-components relations are specified in order to state constraints between different timelines. Such a relations compose an overall *domain theory*. In such a way, a component's timeline can be related with other different components' behavior. Synchronizations between pointing's and uncontrollable timelines constitute the set of domain theory constraints. Science operations occurrences during pericentre orbits or ground station availability needed during communications are examples of such a constraints.

A solution plan is represented as a set of decisions on the components' behaviors. Focusing on a singular component, the solution plan is a sequence of allowed values that the component has to assume in a given time frame. Also the time points, at which such a changes occur, are fixed and provided within the plan.

For each component, a clock and a counter are generated. The clock counts how many times a value holds in a component, while the counter enumerates the plan steps within the behavior.

For each component, a corresponding automaton is generated stating its consistent behaviors. In fact, for each corresponding consistency features an appropriate transition is defined. Thus, temporal constraints can be stated as temporal guards on transitions or as state invariants.

Then, for each timeline, a corresponding automaton is generated and appropriate transitions are introduced to define values changes.

In order to validate plan behaviors, components and planned values are to be related within the time frame considered. That is, components and planned timelines have to follow the same evolution. In this way, it is guaranteed that the solution plan is consistent w.r.t. the domain model. Such a relation, can be defined by synchronizing plan behaviors and components changes. Thus, when a plan behavior requires a certain value for a component, the related component automata has to assume such a value.

Once such a relation is realized, a monitor automaton is generated to check whether plan and domain are consistent or not. If an inconsistency occurs, an error in the plan exists. In such a case the monitor changes its status from *normal* to *error*. Of course, the monitor checks also domain theory violations.

Since our goal is not only to validate solution plans, but also to support users during problems diagnosis, we define the monitor automaton with multiple error states. In this way, users can have additional information about which are the failing components or which domain theory relations are violated. Such a feature can be easily introduced in the validation architecture, introducing suitable transitions from normal to failing states.

Referring to MrSPOCK 's scenario, the monitor automata checks if the planned pointing's timeline is consistent with its allowed temporal evolutions. For instance, *Earth, Slew, Science, Slew, ...* (with appropriate durations) is a consistent behavior, while *Earth, Slew, Science, Earth, ...* is not. Moreover, the monitor verifies that all the timelines are consistent w.r.t. the domain theory constraints: for instance, whenever the *Science* status is present on pointing timeline, the orbit events' timeline has to present *Pericentre* status.

Validation Properties

Once the model previously defined is provided as input to model checkers, a validity plan property has to be defined. So, model checkers can perform verification.

In order to state such a property, we refer to monitor automaton. With respect to the model presented above, the property to be checked is the following: *for each timeline, the last plan step can be reached and the monitor remains normal*. Since, model checkers expect a CTL formula, we use the following:

```
AG (last plan steps can be reached) and (Monitor is normal).
```

Where **AG** means that for **A**ll the possible system evolutions is **G**lobally true that the property is satisfied. Whenever the above formula does not hold, model checkers produce an execution trace of the system witnessing that the monitor automaton reach an *error* state. That is, such a trace represents a system execution showing how the timelines are inconsistent with respect to the components' behaviors. Thus, the reported trace can be used to identify the plan error, or the domain inconsistency, and diagnose from which conditions it origins. In this way, completely specified temporal plans can be validated while very useful information can be provided to users whenever an error occurs.

We performed some preliminary tests in order to verify plan validation architecture performances. We run tests on a linux workstation endowed with a 64-bit AMD Athlon CPU (3.5GHz) and 2GB RAM. We validate plans generated by MrSPOCK ranging within 1 to 10 days of activity, handling from 45 to 335 tasks on all the timelines. The results, depicted in Figure 3, show that UPPAAL performs better than NuSMV (with BMC command).

UPPAAL works on-the-fly, which means that it does not preconstruct a global state graph, or Kripke structure, as a prerequisite for the verification of system properties. Since NuSMV performs such a kind of preconstruction, the experimental results collected do not surprise us.

Although the validation technique presented here is preliminary, we have already obtained some interesting results. For example, our verification tool allowed us to detect and solve an inconsistency in MrSPOCK domain. Namely, we found that MrSpock could generate solutions not consistent with the apocentre-maintenance occurences constraint, which is an implicit requirement (i.e. not represented in the temporal model) for the hybid solver. Using the proposed validation architecture, it was possible to detect the inconsistency and, exploiting the additional information provided by the reported trace, diagnose and fix the error.

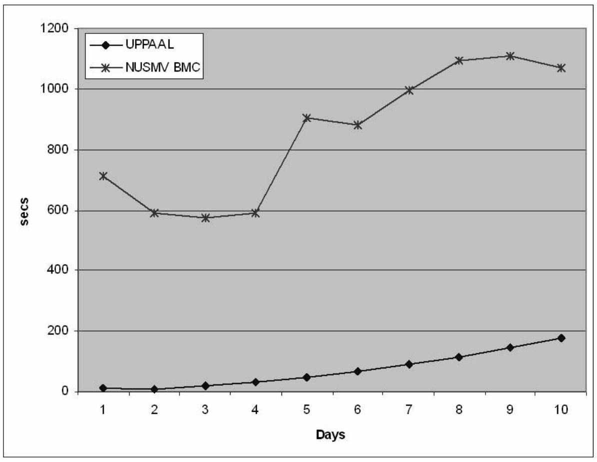

Figure 3: Experimental Results for UPPAAL and NuSMV.

RELATED WORKS

Closely related to our work is the framework proposed in [1]. Here, the authors investigate and compare Constraint Based Temporal Planning techniques and Timed Game Automata methods for representing and solving realistic temporal planning problems. In this direction, they propose a mapping from IxTeT planning problems to UPPAAL-TIGA game-reachability problems and present a comparison of the two planning approaches. This paper is mainly focused on robust plan synthesis while we are interested in plan validation and verification. Indeed, our aim is to address validation and verification issues arising when planning in complex domains and deploying hybrid solving algorithms; these problems are not considered in [1].

Our overall approach is more similar to the one proposed in [19, 9], where model checking (in SPIN) is deployed to guarantee that plans produced by generic automated planning systems meet certain desirable properties. However, differently from our case, the focus is on model validation: planning models are tested and refined to prevent the generation of undesirable plans. The modelling framework and the properties considered in this paper are quite different from the one we are interested in. Indeed, real-time temporal properties and flexible temporally plan verification and validation are not addressed.

A more expressive temporal model is considered in [12] where the authors propose a mapping from interval-based temporal relations models (i.e. DDL models for HSTS) to timed automata models (UPPAAL). This mapping was introduced as a preliminary step towards the application of verification and validation techniques in timeline-based temporal planning, however this direction is not fully developed. In particular, the authors do not propose methods for plan validation.

In the framework of PDDL, we can find the VAL plan validation tool [10] that was extended to permit the validation of plans with durative actions. Although VAL was successfully used during International Planning Competitions since 2002, several open modelling issues are still to be addressed [5].

Formal verification for timeline-based temporal planning is considered also in the ANML framework, a timeline-based specification framework under development at NASA Ames. The work in [17] presents a translator from ANMLite (abstact version of ANML) to the SAL model checker. Given this mapping, the authors present preliminary results to assess the efficency of model checking in plan synthesis. Plan validation and verification issues are not discussed.

A scenario-based verification approach augmented with model-based verification and validation is presented in [11]. The Remote Agent Executor is verified exploiting the Spin model checker. The main focus of the work is on assuring an adequate coverage for scenario-based tests in order to increase confidence on the RAX architecture.

Less closely related, we can find other frameworks where model checking and temporal plan verification are deployed to support robust plan synthesis. For instance, in the CIRCA framework [8], a Controller Synthesis Module (CSM) automatically synthesizes hard real-time reactive plans; the CSM is modeled through a timed automaton model and a model-checking plan verifier is used as a support for robust reactive planning. Here, the main concern is on-the-fly synthesis of control sequences, hence issues and methods (e.g. reactive plan generation and verification) are different from the ones discussed in our paper.

CONCLUSION

In this work, we have presented a preliminary approach to plan validation and verification in a timeline-based planning system. In particular, we have considered V&V issues arising in the MrSPOCK framework, a recent effort for building a timeline-based planning system for the European Space Agency. In this context, plan validation and verification tools are particularly relevant. Indeed, the P&S system of MrSPOCK is based on a hybrid approach where not all the domain constraints can be explicitly represented in the plan domain, therefore the soundness of the plan with respect to the model does not necessarily ensure the soundeness of the produced solution. In this case, an independent solution verifier is needed to test plan consistency with respect to implicit requirements. Furthermore, from the end user perspective, validation and verification tools offer an independent testing environment which can enhance end user trust on the complex and (sometimes) counterintuitive solutions generated by MrSPOCK.

The aim of this paper was to present the approach followed, discussing the problems encountered and the lessons learned in the attempt of deploying V&V techniques and modelling tools in MrSPOCK. In this context, we have presented our integrated plan validation system where plan validation is based on model checking. More specifically, we have defined a translation from our domain models to specifications in UPPAAL and NuSMV in so enabling formal verification of domain and plan properties. In the MrSPOCK domain, we have discussed relevant issues reporting on some preliminary results in temporal plan verification.

A lot of work remains to be done, our long term goal is to obtain a software environment in which V&V technologies are integrated in a complex P&S system and the application developers may take advantage of the co-existence of these tools while knowledge engineering new application.

Acknowledgments. Cesta, Fratini and Tronci are partially supported by the EU project ULISSE (Call "SPA.2007.2.1.01 Space Science"). Orlandini is partially supported by ISTC-CNR. Cesta and Fratini are also partially supported by European Space Agency (ESA) within the Advanced Planning and Scheduling Initiative (APSI).

References

[1] Abdedaim, Y., Asarin, E., Gallien, M., Ingrand, F., Lesire, C., and Sighireanu, M. Planning robust temporal plans: A comparison between cbtp and tga approaches. In *Proc. of ICAPS-2007* (2007), pp. 2–10.

[2] Cesta, A., Cortellessa, G., Fratini, S., and Oddi, A. MrSPOCK: Generating a planning system through a timeline representation framework. In *ASTRA-08. Proceedings of the 10^{th} Workshop on Advanced Space Technologies for Robotics and Automation* (2008), ESA-ESTEC.

[3] Cesta, A., Fratini, S., and Pecora, F. A Multi-Component Framework for Planning and Scheduling Integration. In *PlanSIG'07. Proceedings of the 26th Workshop of the UK Planning and Scheduling Special Interest Group, Prague, Czech Republic, December 17-18* (2007).

[4] Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., and Tacchella, A. Nusmv 2: An opensource tool for symbolic model checking. In *14th Int. Conf. on Computer-Aided Verification (CAV'2002)* (2002).

[5] Fox, M., Long, D., Baldwin, L., Wilson, G., Woods, M., Jameux, D., and Aylett, R. On-board timeline validation and repair: A feasibility study. In *Proceedings of 5th International Workshop on Planning and Scheduling for Space* (2006), M. Giuliano, Ed.

[6] Frank, J., and Jonsson, A. Constraint based attribute and interval planning. *Journal of Constraints 8(4)* (2003), 339–364.

[7] Fratini, S., Pecora, F., and Cesta, A. Unifying planning and scheduling as timelines in a component-based perspective. *Archives of Control Sciences 18*, 2 (2008), 231–271.

[8] Goldman, R. P., Musliner, D. J., and Pelican, M. J. Exploiting implicit representations in timed automaton verification for controller synthesis. In *Proc. of the Hybrid Systems: Computation and Control Workshop* (2002).

[9] Havelund, K., Groce, A., Holzmann, G., Joshi, R., and Smith, M. Automated testing of planning models. In *Proceedings of the Fifth International Workshop on Model Checking and Artificial Intelligence* (2008), pp. 5–17.

[10] Howey, R., and Long, D. Val's progress: The automatic validation tool for pddl2.1 used in the international planning competition. In *Proc. of the ICAPS 2003 workshop on The Competition: Impact, Organization, Evaluation, Benchmarks* (Trento, Italy, June 2003), pp. 28–37.

[11] Jonsson, A., Morris, P., Muscettola, N., Rajan, K., and Smith, B. Planning in Interplanetary Space: Theory and Practice. In *Proc. of the Fifth Int. Conf. on Artificial Intelligence Planning and Scheduling* (2000).

[12] Khatib, L., Muscettola, N., and Havelund, K. Mapping temporal planning constraints into timed automata. In *TIME-01. The Eigth International Symposium on Temporal Representation and Reasoning* (2001), pp. 21–27.

[13] Larsen, K. G., Pettersson, P., and Yi, W. UPPAAL in a nutshell. *International Journal on Software Tools for Technology Transfer 1*, 1-2 (1997), 134–152.

[14] Lowry, M. R., Havelund, K., and Penix, J. Verification and validation of ai systems that control deep-space spacecraft. In *ISMIS-97* (London, UK, 1997), Springer-Verlag, pp. 35–47.

[15] Muscettola, N. HSTS: Integrating Planning and Scheduling. In *Intelligent Scheduling*, Zweben, M. and Fox, M.S., Ed. Morgan Kauffmann, 1994.

[16] Pecheur, C., and Simmons, R. G. From livingstone to smv. In *Proc. of the First International Workshop on Formal Approaches to Agent-Based Systems-Revised Papers* (London, UK, 2001), Springer-Verlag, pp. 103–113.

[17] Siminiceanu, R. I., Butler, R. W., and Munoz, C. A. Experimental evaluation of a planning language suitable for formal verification. In *Proceedings of the Fifth International Workshop on Model Checking and Artificial Intelligence* (2008), pp. 18–34.

[18] Smith, D., Frank, J., and Jonsson, A. Bridging the gap between planning and scheduling. *Knowledge Engineering Review 15*, 1 (2000), 47–83.

[19] Smith, M. H., Holzmann, G. J., Cucullu, G. C., and Smith, B. D. Model checking autonomous planners: Even the best laid plans must be verified. In *Aerospace, 2005 IEEE Conference 5-12 March 2005* (2005), IEEE Computer Society, pp. 1 – 11.