

An Introduction to the J-TRE Environment and its Features

Riccardo De Benedictis

CNR – National Research Council of Italy, ISTC
name.surname@istc.cnr.it

Abstract

Common timeline-based planners are defined as complex software environments suitable for generating planning applications, but quite heavy to foster research work on specific open issues. The J-TRE environment represents a general architecture for timeline-based reasoning that brings together key aspects of such reasoning leaving freedom to specific implementations on both constraint reasoning engines and resolution heuristics. The effectiveness and flexibility of the proposed approach is demonstrated by two real-world applications, based on the J-TRE framework, both in a classical space domain and in a totally novel e-learning domain.

Introduction

Common timeline-based planners like EUROPA (Jonsson et al. 2000), ASPEN (Chien et al. 2010), IxTET (Ghallab and Laruelle 1994) and TRF (Fratini, Pecora, and Cesta 2008; Cesta et al. 2009) are defined as complex software environments suitable for generating planning applications, but quite heavy to foster research work on specific open issues. Some theoretical work on timeline-based planning like (Frank and Jonsson 2003) was mostly dedicated to explain details of (Jonsson et al. 2000) identifying connection with classical planning a-la PDDL (Fox and Long 2003). The work on IxTET and TRF have tried to clarify some key underlying principles but mostly succeeded in underscoring the role of time and resource reasoning (Cesta, A. and Oddi, A. 1996; Laborie 2003). The search control part has always remained significantly under explored. The current realm is that although these planners capture elements that are very relevant for applications, their theories are often quite challenging from a computational point of view and their performance are rather weak compared with those of state of the art classical planners. Indeed, timeline-based planners are mostly based on the notion of partial order planning (Weld 1994) and have almost neglected advantages in classical planning triggered from the use of GRAPHPLAN and/or modern heuristic search (Blum and Furst 1995; Bonet and Geffner 2001). Furthermore, these architectures rely on a clear distinction between temporal reasoning and

other forms of constraint reasoning and there is no sign of attempts to change.

The J-TRE environment represents a general architecture for timeline-based reasoning that brings together key aspects of such reasoning leaving freedom to specific implementations on both constraint reasoning engines and resolution heuristics.

Basics on Timelines

To include time into a logic formalism we choose to provide the predicates with extra arguments belonging to the Time domain \mathbb{T} (real or discrete). For example, a predicate $At(l)$, denoting the fact that an agent is at a certain location l , can be extended with two temporal arguments $s \in \mathbb{T}$ and $e \in \mathbb{T}$, with $s \leq e$, representing its starting and ending times, respectively; the $At(l, s, e)$ formula would be true only if the agent is at location l from time s to time e . Similarly to what described in (Muscettola 1994), we call *token* a proposition that has temporal arguments.

In order to force the proposition arguments to assume the desired values, J-TRE allows the imposition of any kind of linear constraints among the arguments and/or between the arguments and other variables as, for example, quantitative temporal interval relations (Allen 1983). Since constraints must often be customized by the user, the J-TRE framework facilitate the synthesis of planning domains by allowing the organization of constraints in macros called *relations*.

The task of the solver is to find a legal sequence of tokens that bring the timelines (that constitute the partial *plan*) into a final configuration that verifies both the *domain theory*¹ as well as a determined set of desired conditions called *goals*. Starting from an initial state, the planner moves in the partial-plan search space by adding or removing tokens and/or relations (i.e., changing the current state) until all goals are satisfied.

From a planning perspective, the easiest way to describe a *timeline* is to consider it a mere collection of tokens. The predicates that can be accommodated on a timeline as well as the behavior assumed by the planner when a new token is added to a timeline depend on the nature of the timeline itself and, in some cases, on the modeled domain. J-TRE allows the utilization of families of timelines which provide

¹The set of rules that model the domain's dynamic behavior

different modeling ability, such as multi-valued state variables (Muscettola 1994) as well as renewable and consumable resources like those commonly used in constraint-based scheduling (Laborie 2003).

The *state variable* is the most used type of timeline in this approach to planning. State variable predicates are defined by the user during domain definition. The semantics of a state variable is that for each time instant $t \in \mathbb{T}$ the timeline can assume only one value. This corresponds to a mutual exclusion rule between different tokens. Let us assume, for example, to have a predicate $At(l, s, e)$ and a predicate $GoingTo(l, s, e)$. We know for sure that tokens assuming At and $GoingTo$ values cannot overlap. However, two tokens both assuming the At value can overlap if and only if their respective parameters (l , s and e) are pairwise constrained to be equal. In this case we talk about *merging* (or, in some cases, *unification*) of tokens.

Now let us suppose that we want a rule stating that every time we are going to a given location we will reach that location. We can enforce such rule by temporally constraining the $GoingTo(l, s, e)$ and the $At(l, s, e)$ predicates by means of the Allen's *meet* relation having the same location l (see (Allen 1983)). In other words, for each token t_i with a $GoingTo(l, s, e)$ value the planner must ensure that t_i *meets* another token t_j with an $At(l, s, e)$ value, either by imposing the *meet* constraint between t_i and t_j if they both exist, or by adding the missing token t_j before enforcing the same constraints. This kind of "rules" are generalized in a concept usually called *compatibility* (again, here we use a terminology consistent with (Muscettola 1994)). Compatibilities define causal relations that must be complied to in order for a given token to be valid. Although the syntax can be quite different among various planners, a compatibility can be recursively defined by means of a *reference* predicate and a *requirement* where a requirement can be a slave (or target) predicate, a relation among predicates, a conjunction of requirements or, in rare cases, a disjunction of requirements. It is important to underscore that the compatibilities may often involve predicates defined on different timelines, thus allowing to synchronize concurrent activities on different domain components. Most timeline-based planners admit only conjunctions of requirements and reproduce disjunctions by assigning multiple compatibilities to the same predicate.

To simplify matters, we describe compatibilities through logic implications $reference \rightarrow requirement$. In some cases, we will give tokens a specific name and will address their arguments using a Java style *dot* notation (i.e., given a token t having proposition $P(s, e)$ its starting point will be expressed as $t.s$).

Other commonly used types of timelines are the *resources* characterized by a *resource level* $\mathcal{L} : \mathbb{T} \rightarrow \mathbb{R}$, representing the amount of available resource at any given time, and by a *resource capacity* $\mathcal{C} \in \mathbb{R}$, representing the physical limit of the available resource.

We can identify several types of resources depending on how the resource level can be increased or decreased in time. A *consumable resource* is a resource whose level is increased or decreased by some activities in the system. An example of consumable resource is a reservoir which

is produced when a plan activity "fills" it (i.e., a tank refueling task) as well as consumed if a plan activity "empties" it (i.e., driving a car uses gas). We model consumable resources through a timeline that has two allowed predicates: a predicate $produce(a, s, e)$ to represent a resource production of amount a from time s to time e and a predicate $consume(a, s, e)$ to represent a resource consumption of amount a from time s to time e . The planner may need to identify an ordering of the involved activities in order to avoid overproductions (resource level \mathcal{L} cannot exceed capacity \mathcal{C}) as well as overconsumption (resource level \mathcal{L} cannot be lower than zero).

The last commonly used timeline type, quite popular in the scheduling literature, is the *reusable resource*. Reusable resources can be modeled as consumable resources that are produced at their start time and are consumed at their end time. We can model reusable resources through a predicate $use(a, s, e)$ that is true iff there is a production of resource of amount a at time s and a consumption of resource of amount a at time e . Now let's assume we have two tokens t_0 and t_1 belonging to a reusable resource timeline such that $t_0.s < t_1.e \wedge t_1.s < t_0.e$ (this constraint simply forces their overlapping). The expected behavior of the resource is to have a resource usage of $t_0.a$ during t_0 's duration when there isn't overlapping with t_1 , a resource usage of $t_0.a + t_1.a$ when t_0 overlaps with t_1 , a resource usage of $t_1.a$ during t_1 's duration when there is no overlapping with t_0 and a resource usage of 0 elsewhere.

Reasoning on Timelines

Having defined the basic terminology to describe a partial plan and the timelines, we address the problem of reasoning with such building blocks introducing first our planning architecture. A *Domain Definition Language* called DDL.4 is the entry point to the J-TRE environment, allowing the final user to specify the domain objects of interest, the relevant physical constraints that influence their possible temporal evolutions (e.g., compatibility/coordination constraints among different objects, maximum capacity of resources, etc.), as well as the planning goals.

Common timeline-based planners reach a solution state by applying an iterative refinement procedure. If we call *flaw* every possible inconsistency of the current plan, the role of the planner can be reduced to the identification and the resolution of each flaw in the plan. The planning process proceeds until a consistent plan is found, i.e., the propagation of the solving constraints succeeds and all flaws are eliminated. The general solving strategy broadly entails: (i) identifying a set of flaws, (ii) selecting a flaw according to a *selection strategy*, and (iii) solving it by using a *resolution strategy* (see Figure 1 as a reference). During the solving process, a consistency check routine is called on each domain object, possibly generating new flaws to be solved. The solving procedure ends when a consistent node (i.e., containing no flaws) is found.

While flaws can be of different types and can arise for different reasons, what they all have in common is that a *search choice* is necessary to solve each of them. Depending on the reason why a flaw arises, there can basically be four kinds of

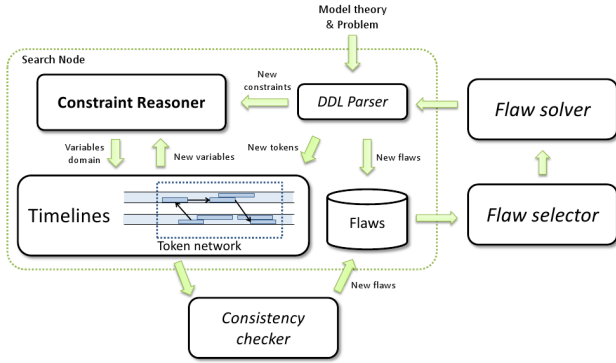


Figure 1: The J-TRE architecture. The planner collects flaws, selects one of them and solves it by executing some DDL code. The planning process will stop in a search space node without flaws.

flaws: (i) *goal flaws*, arising when a new token is added to a state variable to satisfy a compatibility requirement, (ii) *disjunction flaws*, arising when a disjunction statement is found while enforcing some domain rule (expressed in DDL code), (iii) *preference flaws*, arising when a preferred statement is found again while enforcing some domain rule, and (iv) *timeline inconsistency flaws*, arising when an inconsistency is detected on some domain object like, for example, different values overlapping on a state variable, reusable resource oversubscriptions, consumable resources overproductions or overconsumption, etc.

Each time a new node of the search space is created or new constraints are added to the current search space node, a *check consistency* routine is called on each object of the domain and, depending on the object itself, further flaws can be added to the current search space node. This procedure is required in order to remove any further inconsistencies from the timelines *scheduling* tokens in time. This technique has been introduced in (Fratini, Pecora, and Cesta 2008), in which state variables are observed as resources over time and contentions peaks over their continuous representation are removed by adding precedence constraints among tokens.

While, in our system, there is almost no difference in *which* flaw is solved first (as far as we ignore efficiency aspects) because they all have to be solved sooner or later, there could be serious troubles in *how* they are solved, especially in case of cyclic problems.

Consider, for example, a simple state variable having $At(l, s, e)$ and $GoingTo(l, s, e)$ as allowed values. Moreover there is a compatibility for predicate At that requires for each token to start at 0 *or* to be met by a $GoingTo$ token with same location. Finally, a compatibility for predicate $GoingTo$ that requires for each token to be met by a predicate At . We have an initial state with a token $At(l_0, 0, [1, +inf])$ and a goal $At(l_3, [0, +inf], [1, +inf])$. The planner has to apply related compatibility for the goal token producing a sub-goal $GoingTo(l_3, [0, +inf], [1, +inf])$ than another sub-

goal $At(l, [0, +inf], [1, +inf])$ that can unify with first token or apply another compatibility resulting in another $GoingTo(l, [0, +inf], [1, +inf])$ possibly leading to an infinite loop planning about the agent going walking around. In short, although scheduling search space, however exponential, is always finite, it can be the case that compatibility application space is infinite.

Although a crafty strategy does not exist yet (exception made for some work by Bernardini (Bernardini and Smith 2007)) the idea we have pursued is to proceed in depth on the search space maintaining a bound on the number of conflicts. If failing to solve a node within that bound, the overall solving procedure will backtrack to the lowest possible level restarting the search. However, additional solutions to this problem still need to be investigated.

Among the advancements offered by the J-TRE software infrastructure w.r.t. to previous timeline representation frameworks, such as the TRF (Cesta et al. 2009), we underscore the following: (i) the “unification” of the concept of *flaw* (i.e., a plan inconsistency) into a single entity that is uniformly treated (and reasoned upon) throughout the whole J-TRE infrastructure. In J-TRE, flaw analysis and management is no longer spread across specialized reasoners depending on the flaw type, thus allowing to introduce more effective search heuristics that exploit the cross-comparisons among flaws of different types; (ii) the possibility to express constraint of increasing complexity among different domain parameters (e.g., modeling the dependency between resource quantity to be produced and the production activity duration, etc.); (iii) the introduction of the consumable resources among the timeline types.

Current uses of J-TRE

The J-TRE system has been successfully applied in a number of practical applications demonstrating effectiveness and flexibility of the proposed approach. (Cesta et al. 2013), for example, describes the use of the J-TRE environment in a typical space domain where planning is used to organize activities to support facility management on the ISS. An interesting and novel usage of the J-TRE planner is constituted by the PANDORA project (Cortellessa, De Benedictis, and Pagani 2013) and is hinted in the following section. Current work is aimed at developing heuristics needed to build a domain independent planner with performance that are comparable with classical planners.

Continuous Plan Repair in PANDORA

PANDORA is a multimedia training system that utilizes J-TRE as reasoning back-end. Goal of the PANDORA project is to build an intelligent training environment able to deploy a spectrum of realistic simulations of crisis scenarios that: (1) reproduce the stressful factors of the real world crisis; (2) personalize the planned stimuli according to the assessed abilities and features of different trainees and (3) supports the dynamic adaptation of “lesson plans” during the training time-horizon.

A natural technology for achieving these tasks has been identified in the timeline-based planning. The PANDORA

system makes large use of the J-TRE environment to model a number of domain features. Planning is used to compute diversified evolutions of the crisis scenario which correspond to alternative training paths. In addition, planning technology is used to model and maintain trainees’ behavioral patterns according to which aspects of the training can be personalized. The idea of using planning within PANDORA is connected to the synthesis of a “lesson plan”, that is an organized set of lesson’s items, represented through J-TRE tokens, which are given to trainees over a span of time according to a given training strategy. Goal conditions are characterized by high level scenario events representing the abstract blueprint for the master plan. The scenario represents an abstract plan sketch that works as a sequence of “lesson goals” and as a skeleton plan for the ground planner. It is described through a particular “Scenario” timeline that generates sub-goaling by interacting with the set of compatibilities. Disjunctions of requirements produce branches on the search tree guaranteeing varieties of presented scenarios. In particular, it may happen that some compatibility cannot be applied since it imposes too strict constraints resulting in an inconsistent partial plan. In such cases, a the procedure allows to go back to the highest safe decision level.

It is worth saying that the *user’s psychological status* during the training is assessed through psychological self-assessment and physiological measurement, and is then represented by means of similar temporal items so as to insert also these data in a uniform structure and use causal connections between different parts of such plan to foster the continuous update of the plan.

Starting from scenario goals and from the set of domain compatibilities, the planning process generates a plan that is consistent with the given goals, ordering tokens in time through scheduling features and producing proper event consequences. Additionally new goals can be added during crisis simulation to represent (a) decisions taken by trainees, (b) inferences made by the behavioral reasoner, (c) new scenario steps added by the trainer. Within PANDORA, the J-TRE planner is therefore able to replan in order to make its current partial plan to remain consistent with respect to the new dynamic input and with its consequences, namely, changing the current course of the simulated crisis.

Figure 2 shows an example of crisis plan distributed over three timelines representing (a) the number of available police officers, (b) the scenario timeline and (c) actions taken by a trainee. Arrows represent causal links. This means that $mail_0$, $video_{12}$, $question_{21}$ and $document_{15}$ are there *because of* high level goal “Plane crash snippet”. It is worth underscoring how taking dynamically into consideration the answers received by the trainee can cause a dynamic adaptation of the scenario plan with the insertion of the sub-goal $document_{24}$ and an update on the “police officers” resource (the example is quite trivial for the sake of space) that may change the perception of the current crisis scenario either of a single trainee or to the whole class. This is also a way for increasing the workload to a single person or to the whole class.

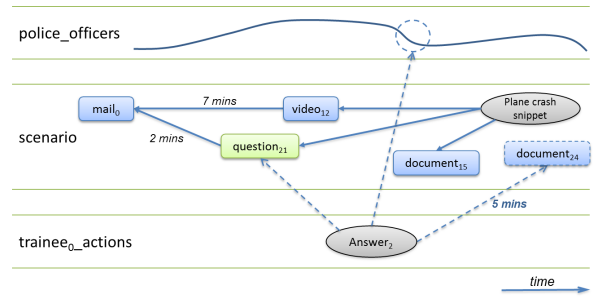


Figure 2: The use of the J-TRE environment for the representation of “lesson plans” in PANDORA.

J-TRE as a Domain Independent Solver

The J-TRE system is currently being tested in several domains. The aim is to gather a large number of problems to be used as benchmarks for domain independent solvers. In the following, a robotic domain extracted from the GOAC project (Ceballos et al. 2011) is introduced. This domain is based on a robotic platform responsible for the movements, a payload represented by stereocameras mounted on a Pan-Tilt Unit, and a communication facility. The planner should arrange activities taking care of on board memory constraints, power requirements and availability of the communication channel.

Figure 3 shows the knowledge contained in the domain by detailing the values that can be assumed by state variables and the legal value transitions in accordance with the mission requirements and the robot physics. To obtain a timeline-based specification of our robotic domain, four state variables have been used, namely, *RobotBase*, *PTU*, *Camera* and *Communication*.

The robot can be in a position ($At(x, y)$) or moving towards a destination ($GoingTo(x, y)$). The PTU can assume a $PointingAt(pan, tilt)$ value if pointing a certain direction, while, when moving, it assumes a $MovingTo(pan, tilt)$ value. The camera either takes a picture of a given object in a position $\langle x, y \rangle$ with the PTU in $\langle pan, tilt \rangle$ ($TakingPicture(f, x, y, pan, tilt)$) or is idle. Similarly, the communication facility can be either operative or dumping a given file ($Communicating(f)$) or idle. The reusable resource *Power* represents consumed power in time while the consumable resource *Memory* represents memory consumption in time. Additionally, one external resource, the *HRDL*, represents contingent events, i.e., the communication channel availability.

During operation, the rover should follow some rules to maintain safe and effective configurations. Namely, the following conditions must hold during the overall mission: (i) while the robot is moving, the PTU must be in the safe position (pan and tilt at 0) and 40W of power are required; (ii) the robotic platform can take a picture only if the robot is still in one of the requested locations while the PTU is pointing to the related direction and if an adequate amount of on board memory is available to store the picture; (iii) once a picture has been taken, the rover has to communicate the picture to the base station; (iv) while communicating the rover has to

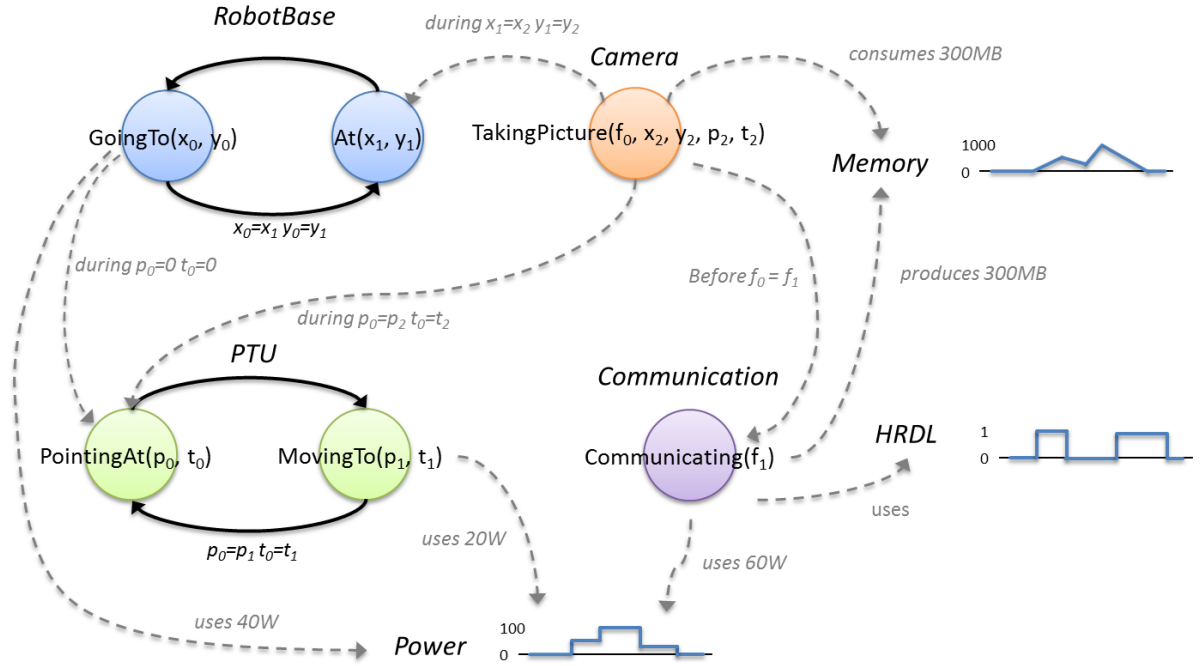


Figure 3: State variables and resources describing the robotic platform and the communication channel availability.

stay still, 60W of power are requested and the memory is released of the amount of transmitted file; (v) while communicating, the orbiter needs to be visible.

As an example, a possible mission action sequence can be the following: navigate to one of the requested locations, move the PTU pointing at the requested direction, take a picture, then, communicate the image to the orbiter during the next available visibility window, put back the PTU in the safe position and, finally, move to the following requested location. Once all the locations have been visited and all the pictures have been communicated, the mission is considered successfully completed.

In (De Benedictis and Cesta 2013) the J-TRE environment has been tested within different variations of the same framework in. By applying the IFPC algorithm (Planken 2008) for incrementally solving temporal problems it has been possible to further improve the performance. Figure 4 shows execution time (in milliseconds) of our benchmark problem, scaled by adding different *TakingPicture* goals (the number of *TakingPicture* goals is on the abscissas) given an initial situation in which the robot is at location $\langle 0, 0 \rangle$ and the pan-tilt is oriented to $\langle 0, 0 \rangle$. We can observe how the J-TRE planner scales quite well with respect to all the problem instances. A comparative evaluation with similar planners (mainly EUROPA, ASPEN, IxTET and TRF) is scheduled.

Conclusions

Temporal flexibility required in controlling mechanisms in real-time (i.e., robotics), interacting with agents requirements as well as uncertainty of real world domains, are just some of the arguments that are leading to the progressive exploration of different planning methodologies and to the extensions of most classic ones.

Timeline based planning constitutes an intuitive alterna-

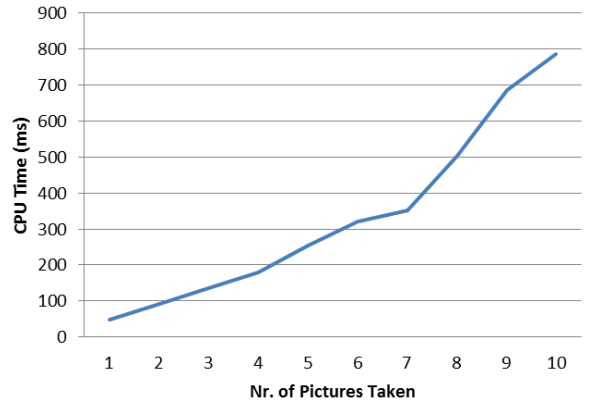


Figure 4: Execution time (in milliseconds) for the GOAC domain (for increasing number of *TakingPicture* goals)

tive to classical planning approaches by identifying relevant domain components evolving in time. Although attractive from a temporal flexibility point of view, these kind of planners have to cope with performance issues due to the complexity that derives from their expressiveness.

The J-TRE architecture filters out common elements from timeline-based planners with the intent to focus the attention to underlying constraint reasoners that could lead to different possible flaw selection and flaw resolution strategies.

Some points are still open. The heuristics for flaw selection and flaw resolution strategies are still relatively poor to compete in performance with state-of-the-art classical planners. The thrust to classical planning given by GRAPHPLAN and the consequent development of heuristic based search is something that is still missing in the timeline-based area.

Acknowledgements The author is thankful to the Planning and Scheduling Technology Laboratory and especially its coordinator Amedeo Cesta.

References

- Allen, J. F. 1983. Maintaining Knowledge about Temporal Intervals. *Commun. ACM* 26(11):832–843.
- Bernardini, S., and Smith, D. 2007. Developing Domain-Independent Search Control for EUROPA2. In *Proceedings of the Workshop on Heuristics for Domain-independent Planning at ICAPS-07*.
- Blum, A., and Furst, M. L. 1995. Fast Planning Through Planning Graph Analysis. In *IJCAI*, 1636–1642. Morgan Kaufmann.
- Bonet, B., and Geffner, H. 2001. Planning as Heuristic Search. *Artificial Intelligence* 129(12):5–33.
- Ceballos, A.; Bensalem, S.; Cesta, A.; de Silva, L.; Fratini, S.; Ingrand, F.; Ocon, J.; Orlandini, A.; Py, F.; Rajan, K.; Rasconi, R.; and van Winnendael, M. 2011. A Goal-oriented Autonomous Controller for Space Exploration. In *ASTRA-11. Proc. 11th Symposium on Advanced Space Technologies in Robotics and Automation*.
- Cesta, A. and Oddi, A. 1996. Gaining Efficiency and Flexibility in the Simple Temporal Problem. In Chittaro, L.; Goodwin, S.; Hamilton, H.; and Montanari, A., eds., *Proceedings of the Third International Workshop on Temporal Representation and Reasoning (TIME-96)*. IEEE Computer Society Press: Los Alamitos, CA.
- Cesta, A.; Cortellessa, G.; Fratini, S.; and Oddi, A. 2009. Developing an End-to-End Planning Application from a Timeline Representation Framework. In *IAAI-09. Proceedings of the 21st Innovative Applications of Artificial Intelligence Conference, Pasadena, CA, USA*.
- Cesta, A.; De Benedictis, R.; Orlandini, A.; Rasconi, R.; Carotenuto, L.; and Ceriello, A. 2013. Integrating Planning and Scheduling in the ISS Fluid Science Laboratory Domain. In *IEA/AIE-2013 26th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems*.
- Chien, S.; Tran, D.; Rabideau, G.; Schaffer, S.; Mandl, D.; and Frye, S. 2010. Timeline-Based Space Operations Scheduling with External Constraints. In *ICAPS-10. Proc. of the 20th Int. Conf. on Automated Planning and Scheduling*.
- Cortellessa, G.; De Benedictis, R.; and Pagani, M. 2013. Timeline-based Planning for Engaging Training Experiences. In *ICAPS-2013 23rd International Conference on Automated Planning and Scheduling*.
- De Benedictis, R., and Cesta, A. 2013. Timeline Planning in the J-TRE Environment. *Agents and Artificial Intelligence. ICAART 2012 Revised Selected Papers*. CCIS(358):218–234.
- Fox, M., and Long, D. 2003. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research* 20:61–124.
- Frank, J., and Jonsson, A. 2003. Constraint-Based Attribute and Interval Planning. *Constraints* 8(4):339–364.
- Fratini, S.; Pecora, F.; and Cesta, A. 2008. Unifying Planning and Scheduling as Timelines in a Component-Based Perspective. *Archives of Control Sciences* 18(2):231–271.
- Ghallab, M., and Laruelle, H. 1994. Representation and Control in IxTeT, a Temporal Planner. In *AIPS-94. Proceedings of the 2nd Int. Conf. on AI Planning and Scheduling*, 61–67.
- Jonsson, A.; Morris, P.; Muscettola, N.; Rajan, K.; and Smith, B. 2000. Planning in Interplanetary Space: Theory and Practice. In *AIPS-00. Proceedings of the Fifth Int. Conf. on AI Planning and Scheduling*.
- Laborie, P. 2003. Algorithms for propagating resource constraints in AI planning and scheduling: existing approaches and new results. *Artificial Intelligence* 143:151–188.
- Muscettola, N. 1994. HSTS: Integrating Planning and Scheduling. In Zweben, M. and Fox, M.S., ed., *Intelligent Scheduling*. Morgan Kauffmann.
- Planken, L. R. 2008. Incrementally Solving the STP by Enforcing Partial Path Consistency. In Aylett, R., and Petillot, I., eds., *Proceedings of the 27th Workshop of the UK Planning and Scheduling Special Interest Group (PlanSIG 2008)*, 87–94.
- Weld, D. S. 1994. An Introduction to Least Commitment Planning. *AI Magazine* 15(4):27–61.