

# Modelling Planning and Scheduling Problems with Time and Resources\*

ROMAN BARTÁK

Department of Theoretical Computer Science and Mathematical Logic

Charles University in Prague

Malostranské náměstí 2/25, 118 00, Praha 1

CZECH REPUBLIC

bartak@kti.mff.cuni.cz <http://kti.mff.cuni.cz/~bartak/>

*Abstract:* Planning and scheduling are close areas but on the other hand, they use very different solving technologies. Recently, there is a strong demand from industry to integrate both approaches into a single system. In the paper we propose a basic framework for modelling planning and scheduling problems that involve reasoning about time and resources. In this framework we go beyond the traditional definitions of planning and scheduling and, from the beginning, we expect integration of both these areas.

*Key-Words:* planning, scheduling, time, resources, modelling

## 1 Introduction

Traditional AI planning tackles the problem of sequencing operators to achieve some goal. In STRIPS-like planning, the operator is defined by pre-conditions and effects, i.e., the pre-conditions must be satisfied to use the operator, and the effects hold after using the operator. The task is to find a sequence of operators starting from a given set of pre-conditions and achieving a given set of effects.

There is no explicit usage of time and resources in traditional planning. In fact, there are no numeric values used so planning methods are based mostly on symbolic manipulation. That is the reason why planning is assumed to be an AI problem rather than a number crunching task. Nevertheless, we can find time and resources behind the traditional planning notions. At least relative time must be assumed if speaking about operator sequencing, i.e., the pre-conditions hold just before we execute the operator and the operator's effect will be true since we execute the operator (until another operator annihilates the effect). Still, traditional planning uses instantaneous operators, i.e., no duration of the operator is assumed. This is OK if we are just sequencing the operators, but, it may cause problems when overlaps of operators are allowed. Moreover, in reality executing the operator takes some time so the planning system should assume this time when looking for a valid sequence of operators. The above observations are reflected in so called durative actions that are included in the recent version of PDDL [9], a modelling language for planning problems, and that are studied in [6].

While time is hidden in semantics of operators, the resources can be encoded in formulas defining pre-conditions and effects. Even one of the earliest planning problems - a block world problem - involved a resource, the robot's hand that moves the blocks over the table. Encoding resource in pre-conditions and effects is a standard way of modelling resources in traditional planning. However, this technique covers only a limited number of resources, we can call them *state resources*. Pre-conditions describe a required state of the resource to execute the operator, e.g., an empty hand, and effects describe a state of the resource after executing the operator, e.g. holding a block.

In reality, the interaction between resources and operators and the integration of time and resources is more complex, e.g. a single resource may execute several operations in parallel. This brings planning to a new level where the quality and feasibility of the plan depends on time and resources too. Planning community is aware of such real-life demand and handling of time and resources is a hot topic in AI planning.

Time and resources play a key role in the areas of scheduling and timetabling too. The scheduling task is to allocate a known set of activities to available resources over time respecting precedence, capacity and other constraints. Timetabling can be seen as a special case of scheduling [20] with different view of space-time (slots) and different objectives. Thus, we will not speak about timetabling separately.

The main difference of scheduling (and timetabling) from planning is that in scheduling we know the structure of activities while planning has to construct this structure. Therefore, when solving real-life problems planning and scheduling modules can be kept separated: first, we plan which activities (operators) are necessary to satisfy the demands and, second, we schedule the

---

\* Research supported by the Grant Agency of the Czech Republic under the contract no. 201/01/0942. An extended version of this paper has been submitted to AIPS 2002 Workshop on Knowledge Engineering Tools and Techniques for AI Planning.

activities to available resources. This could be useful in some problems due to efficiency issues [17] but in other areas, integration of scheduling and planning seems necessary [1, 16]. Note that this integration is not easy because of rather different techniques used to solve problems in planning and scheduling. While planning is based mainly on symbolic manipulation, scheduling uses number crunching techniques from operations research. Recently, constraint satisfaction seems to provide a bridge between these two different technologies so discussions about integration of planning and scheduling are becoming more realistic now. Constraint programming is a widespread technology in scheduling [19]; application of constraint satisfaction techniques to planning problems is described in [3,13,14,18] among others.

When speaking about integration of planning and scheduling, a formal modelling framework to describe such problems is one of the first issues. There exists a de facto standard modelling language PDDL for description of planning problems [10] and this language is being extended to model time [9]. Other approaches in planning attempts to model resources [4,12]. Still, all these approaches have their limitations when describing real-life resources and time.

Surprisingly, there is no system independent language for scheduling problems; at least we are not aware of any such language. There exists a well-known classification of scheduling problems using the triple (machine environment | job characteristics | optimality criterion) by Graham et al. [5]. However, this is an academic classification, not a modelling language to describe a particular problem. Some modelling languages, like STTL [11], exist for timetabling problems but these languages can hardly be extended to general scheduling problems or to planning problems.

In this paper, we describe a framework for integrated description of both planning and scheduling problems. This framework is based on our previous works on modelling scheduling problems enhanced by planning capabilities [1,2] so time and resources play an important role there. We have abstracted from a particular scheduling problem to cover a wider class of problems including pure planning and pure scheduling problems. We concentrate on a basic structure of the framework rather than on particular attributes (even if we mention some attributes to illustrate how the objects are used). This gives us a freedom of designing a generic framework that can be filled by attributes and that way adapted to a particular problem area. This paper is not about solving algorithms, it is merely about the structure of modelling framework covering both planning and scheduling problems.

## 2 Domain modelling

When describing a problem, we can start with the description of the problem area - a domain. This makes the model more general, because it simplifies changes of the model.

What is it a domain? Let us start with a real-life example of industrial scheduling. When scheduling processes in the factory, the problem description consists of the description of the factory, i.e. machines and processes, and the description of demands (orders). In this case, the domain corresponds to the description of the factory and the particular problem consists of the domain and a set of demands. We can say that the domain is a static part of the whole problem that is not changing or the changes are less frequent.

We propose the model for a domain to consist of three basic elements: activities, resources, and recipes. Activity is a basic scheduled/planned object that usually occupies some time and space. Resources define space for processing the activities and recipes describe direct relations between the activities.

### 2.1 Resources

*Resource* is an object that defines space for processing the activity. We will speak about connection between resource and activity later, so let us now concentrate on resource-only features.

Life of the resource, i.e., evolution of the resource in time can be described using a sequence of *states*. For example, the resource oven uses four states load - heat - unload - clean and these states are repeating in a cycle. Some resources, e.g. classroom in timetabling, have only one state. We expect that resource is an object (machine, room etc.) so consumable resources like fuel are modelled using a tank etc. The resource appears in a single state at a given time so the schedule for the resource consists of the sequence of non-overlapping states.

Basically, the model of resource consists of the set of states and transitions among the states (see Figure 1). The transition describes how the resource can change a state. Typically, information about timing is included so we can define minimal and maximal duration of the state, working time for the states, and transition time.

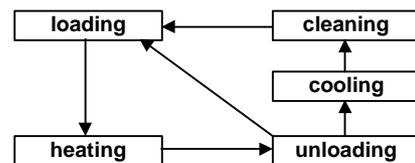


Fig. 1 A state transition diagram for the resource.

Because the resource defines a space for activities, we should describe how much space is available in each state - a state capacity. The state capacity restricts the number of activities that can be processed together. We can also restrict the alignment of activities in the state. Basically, we distinguish between parallel processing, where there is no restriction about the alignment of activities, and batch processing, where the overlapping activities must start and complete at identical times ( see Figure 2).

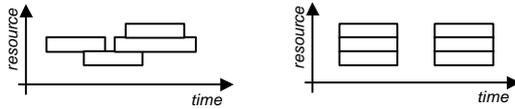


Fig. 2 Parallel (left) vs. batch (right) processing.

To summarise the above discussion, the model of resource consists of the states with some attributes and the transitions between the states (see Figure 3).

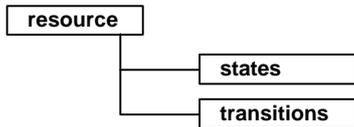


Fig. 3 A basic structure of the resource model.

## 2.2 Activities

*Activity* is a basic scheduled/planned object so when modelling the problem we should specify which activities can be used in the solution. The basic attribute of the activity is its duration, i.e., time occupied by the activity. We can also use time windows to restrict when the activity can be processed.

In many cases, the activity requires some resources for processing. For example, a lecture in timetabling requires a classroom and a teacher, a heating activity in industrial scheduling requires an oven, and a moving activity in transport planning requires fuel. So for each activity we can assign a set of resource requirements. In the *resource requirement* we describe the way of using the resource. Some resources are consumed or produced, we call them consumable resources, and some resources are just used, we call them renewable resources (see Figure 4).

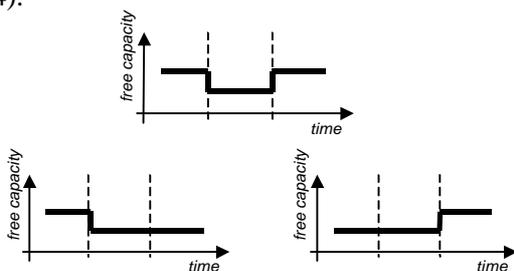


Fig. 4 Renewable (top) and consumable (bottom) resources. Dashed lines indicate start and end of the activity.

Naturally, we should also describe what capacity of the resource is consumed/used/produced. We can also describe what state of the resource the activity requires. Note that the states with batch processing are meaningful for renewable usage of the resource only while parallel processing can be used both for renewable and for consumable usage of the resource.

When specifying the resource requirement, we usually have alternative resources that can satisfy the requirement. Thus we attach a list of resources to each requirement (see Figure 5).

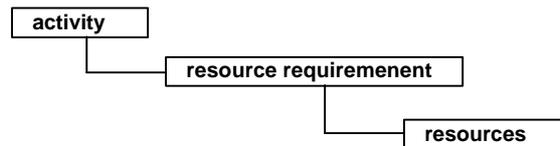


Fig. 5 A basic structure of the activity model.

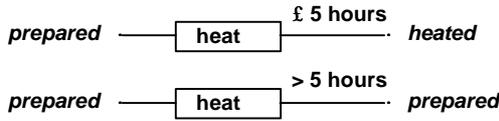
## 2.3 Recipes

The model of activities and resources can describe an indirect relation between the activities only. In particular, the only modelled relation between the activities is via a shared resource, e.g., two activities cannot run in parallel if they share a resource with capacity 1. Such modelling is usually enough for (most) timetabling problems. However, in planning and scheduling we need to model direct relations between the activities (and between the resources), for example a supplier-consumer dependency or a precedence.

Traditional planning uses STRIPS-like rules [8] to model relations between the activities: each activity has some pre-conditions and it generates some effects that may become pre-conditions of another activity. If we add some attributes to the pre-conditions and effects (typically logical terms are used to describe both pre-conditions and effects) we have a general mechanism for information passing between the activities. In HTN (Hierarchical Task Network) Planning [7] the activities are connected into a task graph so more constraints can be expressed over the activities. Moreover, the tasks can be part of another task graph so planning is done via task decomposition and conflict resolution.

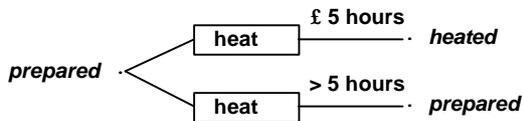
To simplify description of relations between the activities we introduce a notion of *event*. Each activity requires some events to precede it, we say that the activity consumes the events, and each activity generates some other events, we say that the activity produces the events. We call a triple (activity, consumed events, produced events) an *activity environment*. Note that we may have several environments for a single activity, e.g., there exists various combinations of input items consumed by the activity that produces another item (see

Figure 6). Moreover, we can put constraints between the event and the activity, for example to describe the allowed delay between the event and the activity.



**Fig. 6** Two activity environments for a single activity; consumed events are on the left side and produced events are on the right side. Notice also the timing constraint between the activity and the produced event.

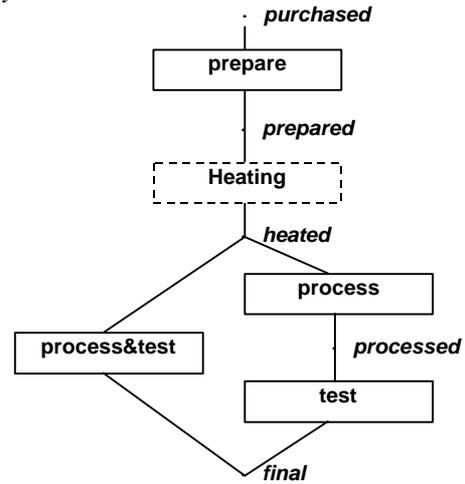
To provide richer modelling capabilities we propose to combine activity environments into a *recipe*. Basically, a recipe is a DAG (directed acyclic graph) where nodes are marked by activities and events. The edge goes either from an activity to an event produced by the activity or the edge goes from an event to the activity that consumes the event. In particular there are no direct edges between the activities and no direct edges between the events. The activity must be connected to all its produced and consumed events (for a given activity environment). So an activity environment forms a sub-graph in the recipe. If there are more environments for the activity then the activity may appear more times in the recipe (each appearance corresponds to one activity environment). However, there are no duplicate events in the recipe. There is one exception when the event may appear two times in the recipe. If the event is produced by one activity and consumed by another activity and connecting both activities to the same event node forms a cycle in the graph. To break the cycle (we require the recipe to be a DAG) we divide the event into two events, one is used as a consumed event only and the other one is used as a produced event only. Let us call such event a *broken event*. Such situation may appear if we want to model recycling or similar features of the real problem (see Figure 7)



**Fig. 7** A primitive recipe Heating; the edges goes from left to right. There is also a broken event "prepared".

In the recipe, there exist three types of events: events that are both produced and consumed (by different activities), events that are produced only, and events that are consumed only. In case of recycling described above, the broken event is part of both consumed-only and produced-only sets of events. Together, the recipe behaves like a meta-activity and thus we can use the recipe within another recipe like an activity environment (see Figure 8).

During planning we are decomposing the required recipes to individual activities but we can also connect different recipes via common events (one recipe produces the event and another recipe consumes the event). Still there could be some events that are consumed only (there is no action that consumes such event); these events may correspond to purchases of raw material etc. Similarly, there could be produced only events, e.g. describing appearance of the final product. We call such produced-only and consumed-only events *one-way events*.



**Fig. 8** A recipe using another recipe (dashed).

If we expect that all the events have unique names then we can represent the recipe as a set of activity environments and recipes (see Figure 9). In such a representation it is clear how the activities and recipes are connected via common events.



**Fig. 9** A basic structure of the recipe.

### 3 Problem modelling

A domain model describes the problem area i.e. which resources are available, what activity types can be used, and what are the relations between the activities. To specify a particular problem we need to describe the actual activities. This could be done explicitly, like in traditional scheduling and timetabling, where the set of activities is given as the input and the task is to allocate the activities to resources respecting the resource and recipe (precedence) constraints. In traditional planning, the input consists of some events and the task is to generate the activities in such a way that the events are connected via activities i.e. the activities in the plan are

described implicitly via the events. In our framework, we propose to combine both these ways of input specification, i.e., depending on the input we will solve either a pure scheduling (timetabling) problem or a pure planning problem or a mixture of both.

### 3.1 Initial data

If we are using resources in the problem, it is a good manner to describe the initial situation/state of each resource. In timetabling this is useless because there are no states. In pure scheduling this is done via specification of the activity with pre-allocation of the activity to the resource and to initial time.

In our framework we allow description of the initial state(s) of each resource as well as specification of activities that are known before we start scheduling. These activities may be pre-allocated, i.e., some of the parameters of the activity are known (like time and used resources) or the parameters are unknown and the task is to find their value (allocate the activity to resources and time). Using such initial data allows us to model pure scheduling and timetabling problems or to use the system to complete partially known schedules. In the second case, new activities are introduced during scheduling to fill gaps in recipes.

### 3.2 Goals

To further extend the planning features of the framework, we allow specification of known events in the description of the problem. Remind that the event make a connection node between the activities. If there appears an event in the system then this event must be produced by some activity and consumed by another activity. Only the one-way events may have either the consumer or the producer. To start planning, we can put some initial events to the system and the system will try to cover them, i.e., to find an action that produces the event and/or the action that consumes the event. Introduction of the action may cause introduction of new events and the task is to cover all the events. As we said above it means that there must be an action producing the event and an action consuming the event. A missing action (producer or consumer) in a one-way event is substituted by including the event among the initial events. Note that this process is similar to STRIPS planning where we have to find activities generating the final effects using the initial pre-conditions.

It is possible that some one-way events are introduced during the process of planning and these events are not included among the initial events. For example we can introduce an event describing a purchase of raw material. To allow such situation we can mark some one-way events as free events. Then, we can

introduce a free event during planning if some activity requires it even if the event is not among the initial events.

To summarise the above paragraphs, the problem is described by specifying the domain (a problem area) and by describing some objects in the final schedule, namely some activities and initial events. The task is to fill the gaps in the schedule following the recipes and respecting the resource constraints (see Figure 10). It means that the resulting plan consists of the activities allocated to resources and connected with other activities via events.

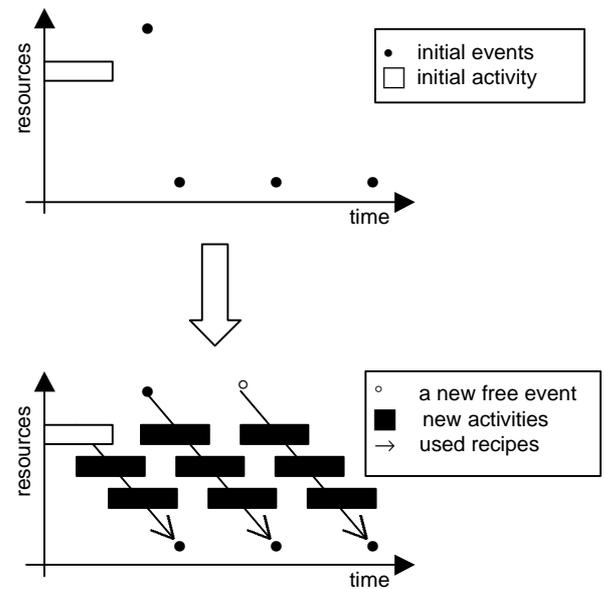


Fig. 10 Gantt charts - from the problem description (top) to the solution (bottom).

## 4 Conclusion

This paper describes the basic concepts behind integrated modelling for planning and scheduling problems. At this stage, we concentrate on the structure of the model rather than on the formal specification of a modelling language that will be the next step. One of the basic requirements in the framework is transparent modelling of time and resources while keeping planning capabilities. We believe that our proposal gives such a transparent modelling framework.

There are several new ideas behind the proposal. First, it is the general relation between the activity and resources which simplifies modelling of alternatives independently on what resource type is used. Second, there is a new concept of events and recipes to describe direct relations between the activities. We took our former models for industrial scheduling problems and we abstracted the item flow into a DAG with events and activities. The idea behind using events is to make the

recipe closer to STRIPS and HTN notions so it is possible to use standard planning techniques for problems described using our framework. Finally, it is the integration issue – we can model pure planning and scheduling problems there as well as we can model problems where planning meets scheduling [1].

*References:*

- [1] Barták R.: On the Boundary of Planning and Scheduling: A Study. In Proceedings of the 18<sup>th</sup> Workshop of the UK Planning and Scheduling SIG, Manchester, pp. 28-39, 1999.
- [2] Barták R. and Rudová H.: Integrated Modelling for Planning, Scheduling, and Timetabling Problems. In Proceedings of PLANSIG 2001, pp. 19-31, Edinburgh, UK, 2001.
- [3] Binh Do M. and Kambhampati S.: Solving planning-graph by compiling it into CSP. In Proceedings of AIPS 2000, pp. 89-91, 2000.
- [4] Brenner M.: A Formal Model for Planning with Time and Resources in Concurrent Domains. In Proceedings of IJCAI-01 Workshop Planning with Resources, Seattle, 2001.
- [5] Brucker P.: Scheduling Algorithms. Springer Verlag, 2001.
- [6] Coddington A., Fox M., Long D.: Handling Durative Actions in Classical Planning Frameworks. In Proceedings of PLANSIG 2001, pp. 44-58, Edinburgh, UK, 2001.
- [7] Erol K., Hendler J., and Nau D.: UMCP: A Sound and Complete Procedure for Hierarchical Task-Network Planning. In Proceedings of 2<sup>nd</sup> International Conference on AI Planning Systems, pp. 249-254, 1994.
- [8] Fikes R. and Nilsson N.J.: STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence 2, pp. 189-208, 1971.
- [9] Fox M. and Long L.: PDDL 2.1: An extension to PDDL for expressing temporal planning domains. Technical Report, Department of Computer Science, University of Durham, UK, 2001.
- [10] Ghallab M., Howe A., Knoblock C., McDermott D., Ram A., Veloso M., Weld D., Wilkins D.: PDDL - The Planning Domain Definition Language, Tech Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.
- [11] Kingston J.H.: Modelling Timetabling Problems with STTL. In Proceedings of The Practice and Theory of Automated Timetabling, LNCS 2079, Springer Verlag, pp. 309-321, 2001.
- [12] Koehler J.: Planning under Resource Constraints. In Proceedings of 13<sup>th</sup> European Conference on Artificial Intelligence, Brighton, pp. 489-493, 1998.
- [13] Laborie P.: Algorithms for Propagating Resource Constraints in AI Planning and Scheduling: Existing Approaches and New Results. In Proceedings of 6<sup>th</sup> European Conference on Planning, Toledo, 2001.
- [14] Nareyek A.: AI Planning in a Constraint Programming Framework. In Proceedings of 3<sup>rd</sup> International Workshop on Communication-Based Systems, 2000.
- [15] Porteous J. and Sebastia L.: Extracting and Ordering Landmarks for Planning. In Proceedings of PLANSIG 2000, pp. 161-174, Milton Keynes, UK, 2000.
- [16] Smith D.E, Frank J., and Jónsson A.K.: Bridging the Gap Between Planning and Scheduling. In Knowledge Engineering Review, 15(1), pp. 61-94, 2000.
- [17] Srivastava B. and Kambhampati S.: Scaling up Planning by teasing out Resource Scheduling. Technical Report ASU CSE TR 99-005, Arizona State University, 1999.
- [18] Van Beek P. and Chen, X.: CPlan: A Constraint Programming Approach to Planning. In Proceedings of AAAI-99, pp. 585-590, 1999.
- [19] Wallace, M.: Applying Constraints for Scheduling, in: Constraint Programming, Mayoh B. and Penjaak J. (Eds.), NATO ASI Series, Springer Verlag, 1994.
- [20] Wren A.: Scheduling, Timetabling and Rostering - A Special Relationship. In Proceedings of The Practice and Theory of Automated Timetabling, LNCS 1153, Springer Verlag, pp. 46-76, 1996.