# On the Boundary of Planning and Scheduling: A Study

Roman Barták[*]

*Charles University, Faculty of Mathematics and Physics*
*Department of Theoretical Computer Science*
*Malostranské námestí 2/25, 118 00 Praha 1, Czech Republic*
*e-mail: bartak@kti.mff.cuni.cz*

---

## Abstract

Planning and scheduling are areas attracting an unceasing attention of computer science community. However, despite of their similar character, planning and scheduling problems are usually handled independently using different methods and technologies. While Artificial Intelligence technology is more relevant for planning problems, Operations Research has a long tradition in studying scheduling problems. Recently, Constraint Programming brings a fresh breeze to both areas as it allows a combination of various methods to solve the planning and scheduling problems within single declarative framework.

In the paper we give a study of mixed planning and scheduling approach to solve the scheduling problems in complex production environments. We analyse three conceptual models and we compare their advantages and limitations from the mixed planning and scheduling point of view. We also give an industrial background to justify the features required in the models. The models were studied within the VisOpt project whose goal is to developed a generic scheduling engine applicable to various complex production environments. However the results can be applied to other (non-production) problem areas where mixed scheduling and planning capabilities are desirable.

*Keywords:* planning, scheduling, mixed approach, modelling

---

## 1 Introduction

Planning and scheduling attract high attention among researches in various areas of computer science. Sometimes, there is confusion what problems planning and scheduling deal with and what are the similarities and the differences.

Roughly speaking, *planning* deals with finding plans to achieve some goal. More precisely, a planning task is defined as finding a sequence of actions that will transfer the initial world into one in which the goal description is true [19]. Naturally, the possible sequences of actions are restricted by constraints describing the limitations of the world. Planning has been studied in Artificial Intelligence (AI) for years and the methods developed there, like the STRIPS representation and planning algorithm [15], are the core of many planning systems.

---

[*] Supported by the Grant Agency of the Czech Republic under the contract no 201/99/D057.

Opposite to planning, *scheduling* deals with the exact allocation of resources to activities over time, i.e., finding a resource that will process the activity and finding the time of processing [9]. Again, the scheduler must respect the precedence, duration, capacity and incompatibility constraints. Operations Research (OR) has a long tradition in studying scheduling problems and many successful methods to deal with the problem were developed there.

In the industry, the border between planning and scheduling tasks is moved to a different level and it becomes a little bit fuzzy. Also, the main difference between traditional planning and scheduling, i.e., the generation of activities in planning vs. assigning activities to resources and time in scheduling, is suppressed here. Both industrial planning and scheduling deal with the task of finding a sequence of activities to achieve some goal and assigning these activities to resources. The main difference is in the resolution of the resulting plan or schedule. While the industrial planning deals with the task of finding “rough” plans for longer period of time where activities are assigned to departments etc., the industrial scheduling deals with the task of finding detail schedules for individual machines for shorter period of time. From this point of view, scheduling can be seen as a high-resolution short-term planning.

The similarity of industrial planning and scheduling brings us to the idea of using a mixed approach that can be applied to both areas. The proposed models will allow a run-time generation of activities (planning) that will be assigned almost immediately to the resources over time (scheduling).

For the problem specification we use the Constraint Programming (CP) that allows declarative description of problems with easily understandable real-life constraints. Moreover, it can encapsulate the solving methods developed in both AI and OR to improve efficiency of the system. *Constraint programming* [7] is based on idea of describing the problem declaratively by means of constraints, logical relations among several unknowns (or variables), and, consequently, finding a solution satisfying all the constraints, i.e., assigning a value to each unknown from respective domain. It is possible to state constraints over various domains, however, currently probably more than 95% of all constraint applications deal with finite domains [22]. And among them, the scheduling problems are the most successful application area [17,23].

In the paper we argue for mixing both planning and scheduling tasks within single system and we study three conceptual models from this point of view. We give a survey of basic techniques, concepts and mechanisms developed within the project of generic scheduling engine. The task was to prepare a generic model capable to capture various planning and scheduling problems in complex production environments. Nevertheless, the results are applicable to non-production areas, like transport problems, as well.

The paper summarises the results of the first stage of the VisOpt project when the basic concepts were studied and the core constraints describing the problem area were introduced. At this stage we concentrated on the expressive power of the models primarily and we only touched slightly the efficiency issues. Next stage of the project will cover the definition of additional, mostly redundant constraints for value propagation as well as a special labelling procedure. At this stage, the efficiency of the system becomes the main issue.

The paper is organised as follows. In Section 2 we specify the problem domain and its specialities. We also highlight the specific problems that require the mixed planning and scheduling approach. Section 3 is dedicated to deep description and comparison of planning and scheduling tasks. We define the traditional planning and scheduling tasks there and we explain their mixture in the industrial planning and scheduling. In Section 4 we describe three conceptual models for complex production environments; we compare their advantages and limits and we list the conditions for successful usage of models. The paper is concluded with some final remarks.

## 2 Problem Area

In the VisOpt scheduling project [6] we deal with complex production areas like plastic, petrochemical, chemical or pharmaceutical industries. The task is to develop a generic scheduling engine that can be customised easily for particular environment via the description of resources, initial situation and expected future situations.

The problem domain can be described as a heterogeneous environment with *several resources* interfering with each other. Currently we are working with producers, movers and stores, later other resources like workers and tools will be added. Some resources can handle several tasks at a time (this is called batch processing) and the task can be scheduled to multiple alternative resources. Also the order of tasks processed by the resource is not arbitrary but the processed task influences what tasks can follow. Consequently, we must follow the *transition patterns* and assume the *set-up times* between the tasks as well. The processing time is usually variable and there is defined a *working time* when the tasks can be processed in resources.

*Alternative processing routes*, *alternative production formulas* and *alternative raw materials* are other typical features of above mentioned industry areas. In addition to the core products it is possible to produce the *by-products*, that can be used as a raw material in further production, or the *co-products*, that can be sold as an alternative to the ordered product. Processing of both by-products and co-products must be scheduled as well because of limited capacity of warehouses where all the products are stored. Also the compatibility constraints describing which products can be processed, i.e., produced, moved or stored together, must be considered. Last but not least there is a possibility of *cycling*, i.e., processing the item for several times for example to change features of the item or just to clean up the store.

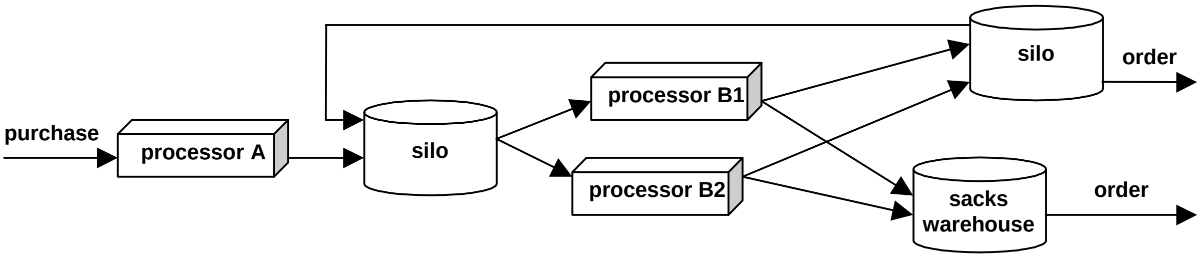

*Figure 1 (example of complex production environment)*

Typically, the production is not driven by the custom orders only but it is possible to schedule the production for store according to the factory patterns and the forecast. It means that a customer ordered not everything that is really produced.

The task is to generate the most profitable schedule or a schedule close to the optimum. Such task is not new for the scheduling community but note that most current scheduling systems use a makespan as the objective function [11,12,13]. The idea of minimising the makespan, i.e., the maximum completion time of the activities, follows the assumption that shorter production time implies lower cost and lower cost implies higher profit. However, this is not necessarily true in many complex production environments where expensive set-ups must be considered. Therefore a more general notion of the cost function should be used as the objective function. Note that we do not describe the cost handling in detail in the paper; the cost parameter is just added to each task and this parameter could be constrained by other task’s parameters like duration. The goal is to minimise the sum of cost parameters for all the tasks in the schedule.

The solved problem is close to the group of resource constrained project scheduling problems (RCPSP). RCPSP [11,14] is a generalisation of job shop scheduling [3] in which tasks can use multiple resources, and resources can have capacity greater than one (more tasks can be processed together). The definition of RCPSP as well as of all other scheduling problems expects the set of tasks to be known before the scheduling starts. Unfortunately, this is not necessarily true in the complex production environments where scheduling the task to a particular resource or time may introduce new tasks to the system. Typically, using alternative processing routes, by-products, co-products and production for store cause such behaviour. Using the foregoing planning phase provides a little help in such cases as we argue in the next chapter.

## 3 Planning and Scheduling: A Comparison

In most current APS (Advanced Planning and Scheduling) systems the planning and scheduling tasks are handled separately in different modules and the communication between the modules is limited. Such decomposition seems natural because the traditional planning and scheduling deal with a bit different tasks and different methods are used to solve the tasks. On the other side, in industry the notions of planning and scheduling are merged and the difference between them is fuzzier.

### *3.1 Traditional Planning and Scheduling*

The traditional definition of planning says that *planning* tackles the problem of finding plans to achieve some goal, i.e., finding a sequence of activities that will transfer the initial world into one in which the goal description is true [19]. It means that a description of the initial world, the specification of the desired world and the list of available activities make the input of the planner. The output consists of the sequence of activities. A typical planning task in the industry consists of finding the production sequences to satisfy the custom orders.

The traditional *scheduling task* deals with the exact allocation of activities to resources (or resources to activities) over time respecting precedence, duration, capacity, and incompatibility constraints [9]. The set of activities, the list of resources and the specification of the constraints make the input to the scheduler. The output of the scheduler consists of the allocation of the activities to the resources over time.

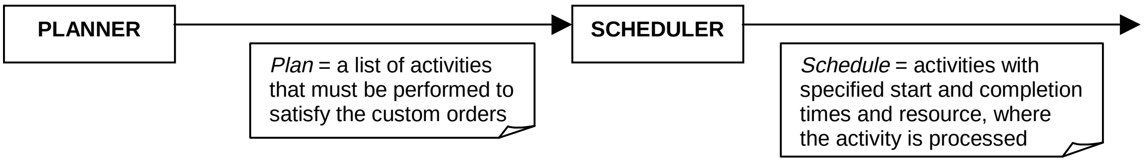

*Plan* = a list of activities that must be performed to satisfy the custom orders

*Schedule* = activities with specified start and completion times and resource, where the activity is processed

*Figure 2 (separate planning and scheduling)*

As Figure 2 shows the communication between separate planning and scheduling modules is simple: the planner prepares the list of activities as well as some constraints, namely the precedence and duration constraints, for the scheduler. The remaining constraints for the scheduling, like the capacity and compatibility constraints, and the list of resources are derived from the factory specification. This simple decomposition is the nicer side of the thing.

On the other side both planner and scheduler should consider the same environment limits and constraints even if the planner uses a more general view while the scheduler deals with the details. In reality the planner does not assume all the production details so it is possible that it generates too tighten plans that are impossible to schedule or it generates plans that are too released, i.e., the plan does not utilise the factory capacity fully. Solving such possibilities requires either more informed planner that must assume the same constraints as the scheduler does or there must be backtrack from the scheduler to the planner to find better plan. This backtrack occurs when the scheduler finds a clash in the plan or it finds that the resources' capacities are not utilised fully. Naturally, backtrack from the scheduling stage to the planning stage complicates the communication between the planning and scheduling modules because the scheduler should inform the planner about the cause of backtrack via identification of the conflict or the not fully utilised resource. Also, the planner must be capable to exploit this additional information so it should handle some scheduling tasks as well. Briefly speaking, the planner should care not only about "what should be processed" but also about "how it should be processed".

Another problem in the traditional definition of the scheduler is that it expects all the activities to be known before the scheduling starts. As we sketched in Chapter 2 such requirement is sometimes inappropriate because the existence of "supplying and consuming" activities may depend on the allocation of the activity to particular resource. The typical example of such behaviour is using alternative processing formulas, alternative raw material or producing by-products. Again, either the planner must consider some scheduling constraints or the scheduler must be able to generate activities during the scheduling that is a typical planning task.

### *3.2 Industrial Planning and Scheduling*

In the real life, the notions of planning and scheduling are not strictly distinguished and sometimes there is confusion between them.

The notion of planning means preparing a plan but what is it a plan? We may have a *marketing plan* that describes the quantities and approximate release times of products using market forecast and current custom orders. This plan is usually for a longer time period and it is more accurate in earlier times than in later times. Notice that the marketing planning has almost nothing in common with the traditional planning described in the previous section. The result of marketing planning consists of the list of demands to the production so there are no sequences of actions that "change the world".

The marketing plan makes the input to *production planning* whose task is to generate a production plan, i.e. a sequence of activities necessary to satisfy the orders (demands) from the marketing plan. The definition of production planning is very close to the traditional planning but the production planning usually covers the allocation of the activities to factory departments as well, that is a typical scheduling task. Production planning uses information like BOM (bill of materials) to generate processing routes and to find what raw material should be ordered and when. Again, the production plan is prepared for a longer period of time.

Finally, there is a *production scheduling* which allocates the activities from the production plan to particular resources over time. The scheduler works with the detail information about the resources, like capacity and compatibility constraints, and the resulting schedule is prepared for a shorter period of time than the production plan (because of efficiency issues and unexpected changes in the environment). The definition of production scheduling is very close to the traditional scheduling, but sometimes during scheduling we need to introduce new activities to process by-products etc.

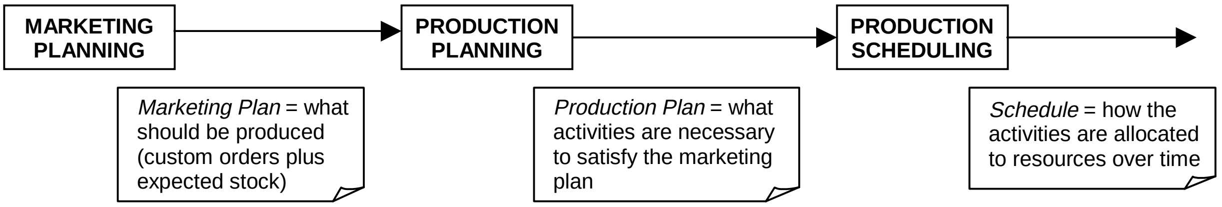

*Figure 3(planning and scheduling in industry)*

As you see, there is no conceptual difference between the production planning and scheduling. Both tasks cover the generating of activities as well as assigning the activities to the resources. The resolution of the result is the main difference between the production planning and scheduling. While the production planning works with departments and a longer period of time, the production scheduling handles individual machines in shorter period of time.

The similarity between the production planning and the production scheduling brings us to the idea of handling both tasks together within single mixed framework.

### *3.3 Mixed Planning and Scheduling Approach*

Let us now return to the specific features of the problem area that are described in Chapter 2 and analyse them from the planning and scheduling point of view.

First, there are alternative processing routes, alternative production formulas and alternative raw materials. The choice of the alternative is part of the planning task but the information necessary for good decision is available at the scheduler level because the decision depends on particular allocation of activities to resources.

Second, there is a production of the by-products and the low-quality products that are produced as “waste” or during the transition between activities. Again, the planner is responsible for generating the activities to process these products but it is the scheduler that decides what and if any by-product appears by assigning the activity to a particular resource. Note that processing of the by-products should be scheduled as well because they may fill the stores otherwise.

Third, there are transition patterns and set-up times that are usually modelled using special transition or set-up activities [18]. Generation of these activities is part of planning task but the existence of the activities depends on the allocation of other activities to resources that is a scheduling task.

Finally, there is a production for store. Normally, the marketing plan should specify the production of items that are not ordered by real customers. Nevertheless, sometimes it is more appropriate to delegate this decision to the scheduler. For example, it could be cheaper to schedule continuous production, i.e., to add new production activities, than stopping the machine.

The discussions in the above paragraphs and sections justify our proposal of mixing the traditional planning and scheduling tasks into single framework. The idea is not very complicated, instead of generating activities in the planner and assigning them to the resources in the scheduler we propose to mix the generation and the scheduling into a single module. We call this mixed module simple a *production scheduler*. We expect to preserve the separate marketing planner that generates the demands for the production but the production scheduler can schedule non-ordered production too[1].

The production scheduler consists of the *activity generator* (former planner) that generates the activities and the *activity allocator* (former scheduler) that allocates the activities to the resources over time (almost) immediately. By attempting to allocate the activity to the resource after its introduction we can detect the clashes sooner as well as we can remove some alternatives via constraint propagation that restricts the domains of activity parameters.

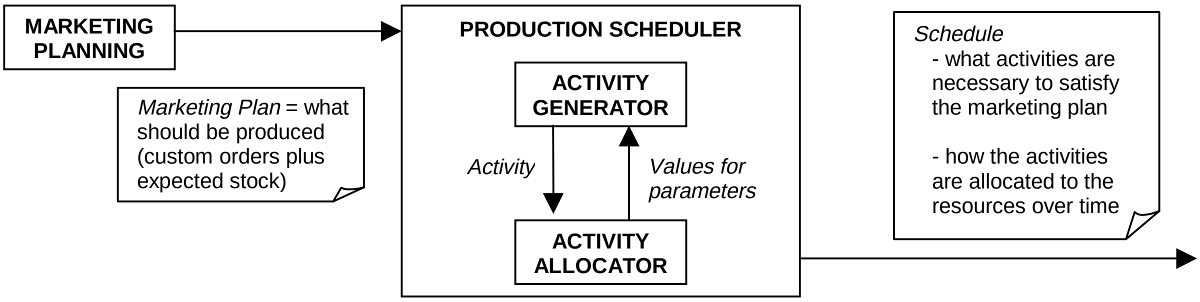

*Figure 4 (mixed planning and scheduling)*

The communication between the generator and the allocator is simple via single activities. The generator introduces activity to the system and asks the allocator to schedule it. The allocator influences the generation of further activities by restricting the domains of activity parameters and it can ask the generator to introduce new activities, e.g., set-ups or transitions. The generator is driven by the set of initial activities that can describe the initial situation as well as the future demands generated by the marketing planner. Also notice that depending on the resolution of the scheduling we can use the production scheduler both for the production planning and for the production scheduling described in the previous chapter.

## 4 Conceptual Models

In the VisOpt project we studied three conceptual models of complex production environments with two different views of time. The time-line model expects time to be discrete, i.e. we are jumping from one time point to another. The second view of time expects event-based time. We consider time steps between interesting events there, in particular between changing activities in the resource. We studied two conceptual models that use event-based time, namely order-centric and resource-centric models that differ in the way of managing dependencies between activities.

### *4.1 Time-Line Model*

The time-line model (also called a timetable approach) is a general method of describing dynamic processes using discrete time intervals. First, we divide the time line into sequence of time slices with identical duration and at each time point (the point between two slices) we describe the situation of each resource using several variables. It is assumed that the behaviour of resource is homogeneous between two consecutive time points, i.e., the key events like changing activity occur only at the edge of two consecutive time slices.

The duration of time slices must be defined according to the duration of activities that can be processed by the resources so it should be a common divisor of activities' duration. Naturally

---

[1] The quantity of non-ordered production is restricted by the cost factor, e.g., assuming the cost of storing.

we prefer longer duration of the slice because it means smaller number of variables and consequently less work to do when the variables are labelled. Together, the *duration of slice* is computed as a greatest common divisor of duration of all activities in all resources.

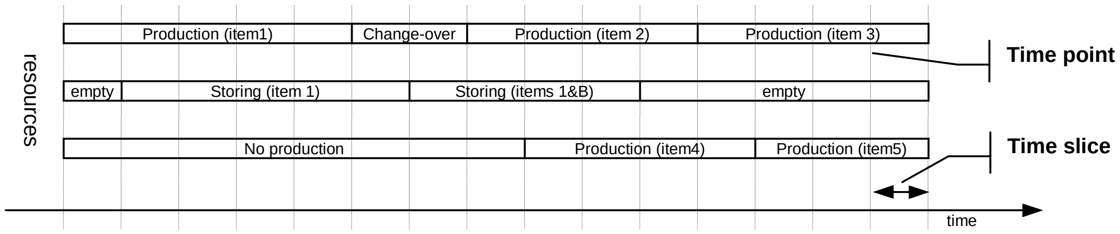

*Figure 5 (a time-line model)*

Now we can describe the situation at each time point using a set of variables. For example in case of store there is a variable for each item that can be stored and this variable specifies the stored quantity. Other variables can specify the state of resource etc. The resource constraints bind variables in the time point and they can express compatibility between stored and processed items or capacity limits of the resource. Also we need constraints binding variables from different time points and different resources. Such constraints express supplier/consumer dependencies or transition patterns. We gave examples of these constraints in [4,5].

Because we know the number of time points in advance (it is derived from the scheduled duration) we can represent the model using a table of variables where one axis corresponds to the variables describing the resources and the second axis corresponds to the time points. This representation requires normalising the resource parameters in all time points.

| ***resource*** | variable name | variables | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| ***Processor A*** | state | | | | | | | | |
| | … | | | | | | | | |
| ***Processor B*** | state | | | | | | | | |
| | … | | | | | | | | |
| ***Store*** | Item 1 quantity | | | | | | | | |
| | Item 2 quantity | | | | | | | | |
| | … | | | | | | | | |
| | **time points** | 1 | 2 | … | | | | | N |

*Figure 6 (a table representation of the time-line model)*

By setting the value of some variables or restricting their domains we can express both the initial situation and the desired future situations that are derived from the marketing plan. Notice that in this model the planning and scheduling components are mixed even more and there is no strict border between them. We don’t need to generate activities because the processed activity in given time is described by the variable in the time point. Consequently, we can use the same constraint propagation technology to solve both the planning and the scheduling task.

Unfortunately the time-line model has the disadvantage of using too many variables when applied to a real-life problem. Even if the activities have longer duration like 25 and 26 minutes we must use one-minute time slice (the common divisor of 25 and 26). Therefore, we can expect not very good efficiency from the time-line model if applied to large-scale scheduling problems. Nevertheless, we believe that the model can be applied successfully in cases when:

- the description of resources is not very complicated, i.e., we use small number of variables to describe the situation,
- the ratio between the time slice duration and the scheduled duration is higher, i.e., either the scheduled duration is not very long or the duration of time slice can be longer.

We also think that the time-line model can be mixed with the event-based models presented in next chapters, in particular to model the store where cumulative scheduling is required.

### 4.2 Order-centric Model

Order-centric model[2] is a traditional model for job-shop scheduling [3] where event-based time is used (event = changing activity in the resource). It is based on idea of defining the chain of activities necessary to produce the ordered item or, generally, to satisfy the order. The goal is to schedule such production chains for all orders respecting the resource limits. In Figure 7 we show an example of production chain. You may notice that by the notion of production chain we mean not only a linear sequence of activities but also, for example, a tree of activities with the root corresponding to the final product or the order.

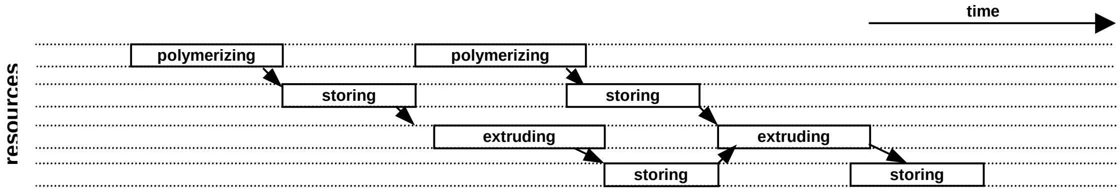

*Figure 7 (a production chain in the order-centric model)*

We describe the activities using a set of variables, typically including the start and completion times and the resource variable. Again there are the resource constraints that bind variables from the activities allocated to the same resource and the supplier/consumer constraints that bind variables from activities in single production chain.

The order-centric model can be represented as a set of production chains where each production chain is represented as a list of activities. The static representation expects all the production chains to be defined before the scheduling starts that complicates the situation if the alternative processing routes, the transition patterns, the set-up activities or the by-products are included. The problem can be solved by introducing virtual activities [18] that are filled during scheduling (by assigning the value to the activity variable) but this complicates the constraints that must include the new activity variable.

The dynamic representation is more appropriate in such case because we can expect that only the first (last) activity in the production chain is initialised, e.g., using the marketing plan, and all other activities are generated by the planning component during scheduling. This solves the problem with alternative processing routes and set-up activities naturally. Handling by-products is more complicated as it requires two production chains to share the activity [4].

Unfortunately, neither the static nor the dynamic representation can solve easily the problems with production for store. This requires the planner to be able to find gaps in the schedule and to introduce a new production chain filling the gaps. Notice that the production for store is different from the processing of the by-products because the processing of the by-products is initialised by one of the original production chains.

### 4.3 Resource-centric Model

A resource-centric model is similar to the order-centric model in the way of using the activities and event-based time. Now, we are working with the list of activities per single resource and the chain of activities per order is handled implicitly by means of dependencies between resources.

---

[2] Sometimes, the model is called task-based model [9] but we prefer the notion of order-centric model as we assign a production chain to the order.

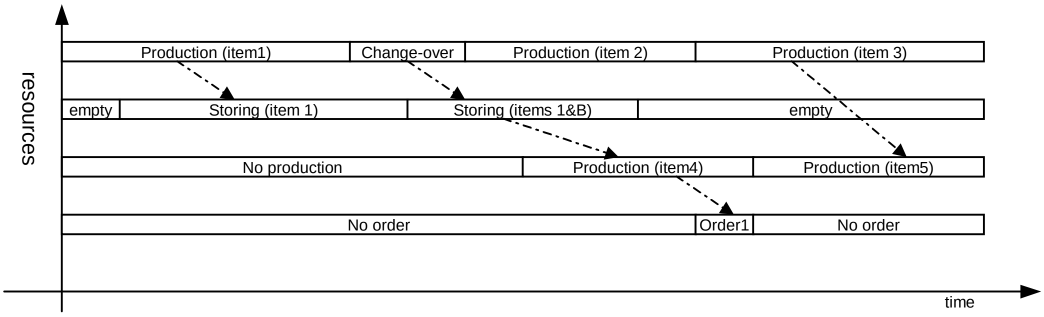

*Figure 8 (a resource-centric model)*

Each resource in the model is described using the set of activity types[3] that can be processed by the resource and using resource constraints. The parameters of the activity (the activity type) are the same as in the order-centric model but we do not need the parameter identifying the resource for activity (each activity is assigned to the resource implicitly) and we need new parameters identifying supplying and consuming activities.

The resource-centric model can be represented as a set of resources where each resource is described by the list of activities. The static representation expects that this data structure is defined in advance, namely we must know the maximal number of activities in the resource (it is derived from the activity duration and the scheduled period). Again we use virtual activities and the scheduler chooses the real activity by assigning the value to the activity variable. In the dynamic model the activities are generated during scheduling for example using the transition schema.

The resource-centric model is able to solve all the problems described in Chapter 2 including the alternative processing routes, the by-products or the production for store. This is because this model describes the factory rather than the production chains that emerge implicitly using the supplier/consumer constraints. In comparison to the order-centric model, it is easier to express various resource constraints there but the planner must identify the “connected” activities to introduce the supplier/consumer constraints [4].

## 5 Concluding Remarks

In the paper we propose a mixed planning and scheduling approach that is capable to solve problems in complex production environments. We also describe three conceptual models to capture such environments and outline their static and dynamic representations (for details see [4,5]).

The static representation requires all the entities and constraints to be introduced before the scheduling starts so we can concentrate on the constraint propagation and labelling only. Notice that it completely removes the separate planning component in cases like alternative processing routes because the scheduler handles everything. On the other hand, the static representation complicates significantly the constraints that must include more parameters. In fact we are using “meta-constraints” consisting of several constraints, one of them is fired when the value of some of the parameters, e.g., the activity parameter, is assigned.

In the dynamic representation we can introduce new entities and constraints during scheduling. In particular we do not need virtual activities because the activity is generated when necessary. Also the constraints are more natural because they can be introduced when

---

[3] Because we describe the abstract activity we use the notion of activity type. During scheduling we use activity types to generate real activities. The relation between activity type and activity is the same as the relation between class and object in OOP.

needed, e.g., when a new activity appears. Also, the dynamic models correspond better to the proposed mixed planning and scheduling approach, where the planner is responsible for the generation of the activities and firing the constraints while the scheduler assigns values to variables, i.e., it allocates the activities precisely.

In the time-line model the differences between the static and the dynamic representation are not significant, mainly because of the static nature of the model (we know the number of time points in advance). The only difference is in the expression of constraints that are more complicated in the static representation (they must include the activity variable). In the dynamic representation the constraints are generated on demand, e.g. when the activity is known.

The order-centric model is a traditional model of job-shop scheduling when order-driven production is used. Usually the static representation is used for this model but if alternatives should be modelled then we propose the dynamic representation that is more transparent. Nevertheless, the order-centric model is still not capable to exploit fully the planning capabilities and the production for the store cannot be modelled.

The resource-centric model is probably the best model for mixed planning and scheduling. It has the same power as the time-line model, i.e. it is capable to capture all the situations as described in Chapter 2, but it removes the main drawback of the time-line model, namely the unnecessary large number of variables. We chose this model as the most promising one for the next research that will cover the search procedures both for the planning component of the systems (what activities should be introduced) and for the scheduler (what values should be assigned to variables).

## 6 Acknowledgements

Author’s work is supported by the Grant Agency of the Czech Republic under contract number 201/99/D057 and by InSol Ltd. I would like to thank Yossi Rissin and the team of InSol for introducing me to the problem and for interesting and encouraging discussions concerning real-life problems of industrial planning and scheduling.

## 7 References

[1] Baptiste, P., Le Pape, C.: A Theoretical and Experimental Comparison of Constraint Propagation Techniques for Disjunctive Scheduling, in: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, Montreal, Canada, 1995

[2] Baptiste, P., Le Pape, C.: Disjunctive Constraints for Manufacturing Scheduling: Principles and Extensions, in: Proceedings of the Third International Conference on Computer Integrated Manufacturing, Singapore, 1995

[3] Baptiste, P., Le Pape, C., Nuijten, W.: Constraint Based Optimisation and Approximation for Job-Shop Scheduling, in Proceedings of the AAAI-SIGMAN Workshop on Intelligent Manufacturing Systems, IJCAIs-95, Montreal, Canada, 1995

[4] Barták, R.: Dynamic Constraint Models for Complex Production Environments, in Proceedings of the 1999 ERCIM/CompulogNet Workshop on Constraints, Paphos, Cyprus, October 1999

[5] Barták, R.: Conceptual Models for Combined Planning and Scheduling, in Proceedings of the CP99 Post-Conference Workshop on Large Scale Combinatorial Optimisation and Constraints, Alexandria, USA, October 1999

[6] Barták, R.: VisOpt – The Solver behind the User Interaction, White Paper, InSol Ltd., Israel, May 1999

[7] Barták, R.: On-line Guide to Constraint Programming, http://kti.mff.cuni.cz/~bartak/constraints/

[8] Bosi, F., Milano, M.: Enhancing CLP Branch and Bound Techniques for Scheduling Problems, Tech. Report DEIS-LIA-98-002, Università di Bologna, 1998

[9] Brusoni, V., Console, L., Lamma, E., Mello, P., Milano, M., Terenziani, P.: Resource-based vs. Task-based Approaches for Scheduling Problems, in: Proceedings of the 9th ISMIS96, LNCS Series, Springer Verlag

[10] Buzzi, S., Lamma, E., Mello, P., Milano, M.: Consistent Orderings for Constraint Satisfaction Scheduling, Tech. Report DEIS-LIA-97-001, Università di Bologna, 1997

[11] Caseau, Y., Laburthe, F.: A Constraint based approach to the RCPSP, in: Proceedings of the CP97 Workshop on Industrial Constraint-Directed Scheduling, Schloss Hagenberg, Austria, November 1997

[12] Caseau, Y., Laburthe, F.: Improved CLP Scheduling with Task Intervals, in: Proceedings of ICLP94, pp. 369-383, MIT Press, 1994

[13] Caseau, Y., Laburthe, F.: Cumulative Scheduling with Task Intervals, in: Proceedings of JICSLP96, pp. 363-377, MIT Press, 1996

[14] Crawford, J.M.: An Approach to Resource Constrained Project Scheduling, in: Artificial Intelligence and Manufacturing Research Planning Workshop, 1996

[15] Fikes, R. E., Nilsson, N. J.: STRIPS: A new approach to the application of theorem proving to problem solving, in: Artificial Intelligence Vol.  No. 3-4, pp. 189-208

[16] Lamma, E., Mello, P., Milano, M., Temporal Constraint Handling in Scheduling Problems, Invited Paper at Intersymp95, Baden-Baden, August 1995

[17] Lever, J., Wallace, M., Richards, B.: Constraint Logic Programming for Scheduling and Planning, in BT Technical Journal, Vol. 13 No. 1, pp. 73-81, 1995

[18] Pegman, M.: Short Term Liquid Metal Scheduling, in: Proceedings of PAPPACT98 Conference, London, 1998

[19] Pool, D., Mackworth, A., Goebel, R.: Computational Intelligence – A Logical Approach, Oxford University Press, Oxford, 1998

[20] Simonis, H., Cornelissens, T.: Modelling Producer/Consumer Constraints, in: Proceedings of CP95, pp. 449-462, 1995

[21] Smith, A.W., Smith, B.M.: Constraint Programming Approaches to Scheduling Problem in Steelmaking, in Proceedings of CP97 Workshop on Industrial Constraint-Directed Scheduling, Schloss Hagenberg, Austria, November 1997

[22] Tsang, E.: Foundations of Constraint Satisfaction, Academic Press, London, 1995

[23] Wallace, M.: Applying Constraints for Scheduling, in: Constraint Programming, Mayoh B. and Penjaak J. (Eds.), NATO ASI Series, Springer Verlag, 1994