Process Design for Efficient Scheduling*

Evgeny Selensky[1], James Little[1], and J. Christopher Beck[2]

[1] Cork Constraint Computation Centre
University College Cork
Cork, Ireland
{e.selensky,j.little}@4c.ucc.ie
[2] Toronto Intelligent Decision Engineering Lab
Mechanical & Industrial Engineering
University of Toronto, Canada
jcb@mie.utoronto.ca

Abstract. In manufacturing, different process designs give rise to different schedules and with each an associated cost. In this paper, we report on a real-life example where a manufacturing company wants to evaluate the scheduling implications related to the degree of coupling between their processes of moulding and casting, in terms of the amount of buffer stock held. The results show that the present configuration could be improved as regards the amount of stock, while still meeting the demand levels. We show this as one example of a process design evaluation and propose in this paper an architecture for generic process design for this company, in order to evaluate quickly other scenarios. From this, we will be able to develop an approach of proactively using scheduling information in a systematic way to positively influence design decisions.

1 Introduction

Manufacturing environments are never static: technology becomes obsolete, demands fluctuate, rapidly changing markets dictate the necessity to vary the range of products, reduce inventory or deploy new resources and so on. Indeed, agility [7] and adaptability [5] of a manufacturing enterprise are increasingly pointed to as a key ingredient for long term economic success. A critical component for making agility and adaptability a reality, is the ability to quickly assess how changes to production processes, factory design and scheduling policies influence the efficiency of their production processes. A company needs to be able to redesign its processes to try to match these new requirements in the environment, but it wants to have some insight into the impact of changes before they are put in place. This introduces serious practical challenges to the company, notably in production planning and scheduling.

Although there has been active research in the area of product design [9], little or no published work has been carried out in how process design influences

* This work has received support from Science Foundation Ireland (Grant 00/PI.1/CO75), Enterprise Ireland Innovation Partnership (IP/2003/160) and Bausch and Lomb.

schedules. A design may introduce several constraints which until the scheduling process happens, do not become apparent. Some previous research [8] has addressed part of this question of the linkage between process design and scheduling. The approach of [8] was such that if the scheduling problem is too difficult, information about resource bottlenecks that often contribute to scheduling difficulty, is made available to the process planner. As a result, process plans can be dynamically re-designed. Although that system allows the factory to react to particular load characteristics of the current orders, it is not capable of reasoning about large-scale resource changes such as machine placement or existence of storage facilities. Our long-term goal is to examine the broader strategic question of making decisions that cannot be reactively modified, for example, due to costs of stopping production to reconfigure the physical layout of the factory or for reasons of similar nature. This paper is a first step in this direction.

We focus on a particular example of a set of changes that may be evaluated, within an architecture that allows a manufacturer to realise scheduling insights of their design choices. We then describe an implementation of this architecture and present results of a preliminary empirical study.

The example is provided by a manufacturer of optical devices, who wants to evaluate the scheduling implications related to the degree of coupling, in terms of inventory held between the processes of moulding and casting. Multiple scenarios needed to be explored, with different demand data sets. Such an evaluation would require considerable time and scheduler resources to determine which is the best configuration. Good quality schedules that are obtained and described in this paper are subsequently evaluated by the user in terms of the company's key performance indicators. The indicators for this company are risk of low stock, satisfying orders within time, resource utilisation and inventory costs.

The paper is organised as follows; section 2 provides a description of the components making up the company's manufacturing process. In section 3, the different coupling scenarios are discussed. In section 4, we describe the models. Section 5 introduces a general architecture according to which we implement and solve the coupling problems. Section 6 then describes the experiments and presents the results. Finally, section 7 derives conclusions and looks towards further extensions within the architecture.

2 Manufacturing Process

The manufacturing process is aimed at producing optical devices of known types in given quantities by the specified due dates. It includes two basic technological steps; moulding followed by casting. Other stages such as inspection and packaging are also present, but not considered significant to the scheduling. Between these two processes there is a store holding completed moulds. Depending on the level of content, this can allow casting to commence immediately without waiting for moulds to be produced. How these two processes are carried out and the intervening stock levels, determine the overall schedule and utilisation of the moulding and casting machines. The various components of the manufacturing process are described below.

2.1 Orders

Based on dynamic market analysis, the company derives quantities and types of product that need to be produced over the next planning horizon. This information is summarised in the form of orders.

2.2 Moulding

At the moulding stage, pairs of moulds are produced that determine the shape of the product. Moulds of different types are made on mould injection machines that operate in cycles. Each machine has a number of cavities. A cavity has its own tooling that needs to be adjusted each time a mould of a different type is produced on it. During such a changeover, all cavities on the machine halt. Often though, the order level implies continuous production of one type of mould throughout the planning horizon. The cycle time and changeover durations are known in advance.

2.3 Casting

Casting is done using specialised casting machines. During a casting operation two different types of mould are clamped together. Plastic is then injected between the moulds to produce a working device.

There is a technological delay between moulding and casting operations to account for mould stabilisation.

2.4 Mould Store

Moulds are stored to minimise the risk of not being able to produce moulds in time for casting and, consequently, of falling behind the due dates. The major sources of risk are the uncertainty in demands and faulty raw material revealed only at testing, after the casting process.

The inventory levels can be kept so as to continuously provide for casting for a certain number of days. In the worst case, when moulding is not possible for some unexpected reason, there always is a period of time that casting can still be done by consuming existing mould inventory while problems with moulding can be resolved. The main disadvantage of this approach is the cost of maintaining the high inventory levels.

In Fig. 1, the possible linkage between the moulding and casting processes is shown.

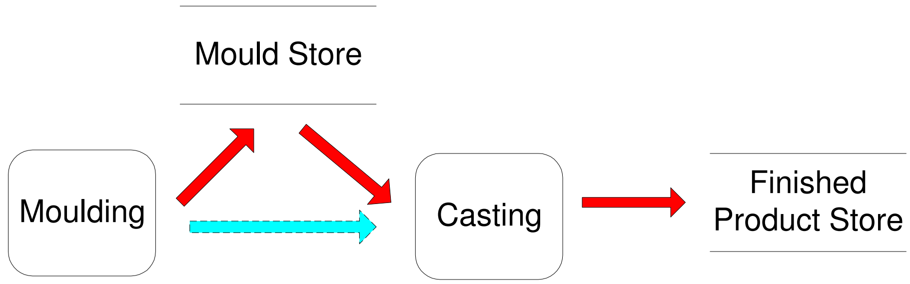

Fig. 1. *Alternative Manufacturing Process Flows.*

3 The Coupling Problem

The company is interested in the degrees of coupling between moulding and casting, to evaluate acceptable trade-offs between stock levels and schedule times. Coupled processes are perceived to involve higher risks as they imply having less mould inventory to react to unpredicted critical changes in the manufacturing environment.

The evaluation of the schedule by the planners is primarily interested in finding solutions that meet the weekly demand. At the same time, the company is interested in other key indicators. Identifying and agreeing a single characteristic to optimize with the company is difficult. Instead, we optimise on makespan and from that, the planners can judge the quality of the schedule through their key performance indicators, described previously.

The coupling scenarios being considered are as follows.

3.1 Scenario 1: Fully Decoupled Process

In this scenario, the company maintains high stock levels sufficient to cover all demand within the planning horizon for immediate casting. Casting is temporally independent of moulding, while the inventory is not exhausted. Casting consumes stock, while moulding, in this case, is performed only to replenish stock. This scenario represents the existing situation on the shop floor and is characterised by a low risk of not meeting the due dates, but high associated storage costs.

3.2 Scenario 2: Fully Coupled Process

In this scenario, no stock is held between the moulding and casting processes. All moulding is performed directly for casting. The moulds are kept in a temporary buffer for a minimum period to attain stability and then supplied for casting. No storage facilities are provided for moulds in this case, they are simply considered as work in progress.

3.3 Scenario 3: Partially Coupled Process

In this scenario, the company holds some stock, but it is not enough to satisfy all orders. There are three production flows, one from stock to casting, one from moulding directly to casting and one from moulding to replenish stock. The amount of moulds that need to be produced directly for casting is assumed to be the difference of production volume and inventory levels over the planning horizon.

4 Manufacturing Process Modelling

We present two models for the process coupling problem. One model is based on a Constraint Programming approach while the other, on an Integer Programming approach. Where the problem can be modelled as a set of linear constraints we can use both approaches, otherwise CP is most convenient.

CP and IP are acknowledged to be competitive on certain problem classes [3]. It turns out that for the decoupled model it is possible to obtain a linear integer formulation and thereby use an Integer Programming (IP) solver, such as Ilog CPLEX [2, 6]. The decoupled model presents no temporal constraints between the moulding and casting and so the problem can be decomposed and solved separately.

4.1 Constraint-Based Models

The constraint-based approach to scheduling uses the notions of resources, activities and temporal constraints as the modelling components [1]. Resources have a capacity and represent either machines on the shop floor or inventory. An activity is a manufacturing operation that has a start time and duration, and that may or may not use resources. Activities are associated with the moulding and casting processes. The number of activities is determined in advance based on the maximum batch size. Unless we are considering a decoupled process, moulding is followed by casting.

Each moulding activity requires exclusive use of a cavity on a moulding machine for its entire duration. Likewise, each casting activity requires exclusive use of a casting machine. Moulding cavities and casting machines are considered individual resources of unary capacity that can be either idle or busy processing an activity. Inventory of moulds and that of finished products are also modelled as separate resources. The cavity or machine allocation is generally not known in advance, neither is the time at which each activity takes place. This is determined during the search. The model is represented in a CP modelling language as follows:

$$\begin{array}{lll} moulding[i] & \underline{requires} & (1)\ MouldingCavities \\ casting[i] & \underline{requires} & (1)\ CastingMachines \end{array}$$

where i is the activity index.

Casting must follow moulding after a sufficient technical delay and each pair of moulds for casting are complete:

$$moulding\_pair1[i].end + Delay \leq casting[i].start$$
$$moulding\_pair2[i].end + Delay \leq casting[i].start$$

Moulding produces stock:

$$moulding[i]\ \underline{produces}\ (amount[i])\ MouldStock.$$

When casting commences, it removes mould stock and, on completion, puts finished product into the finished product store:

$$casting[i]\ \underline{requires}\ (amount[i])\ MouldStock$$
$$casting[i]\ \underline{produces}\ (amount[i])\ ProductStock$$

A special purpose search is adopted which has been found in practice to find good solutions quickly, with the data sets we have been using. This dynamic strategy consists in allocating resources to activities, and setting start times for the activities as follows. It first allocates resources to activities, ordered in decreasing duration, in a round-robin fashion. It then selects the earliest start time t of all activities and chooses an activity that can be scheduled at time t. After that, it considers two alternatives: to schedule the activity at time t or to postpone this activity. This is done for all activities that are not yet scheduled or postponed. A postponed activity is reconsidered every time its start time is updated during search.

4.2 IP-Based Model

The IP models for moulding and casting are similar. We model the manufacturing process as a variant of the classical bin packing problem [4], where it is necessary to determine how to put the most objects in the least number of fixed space bins. In the case of moulding, each cavity is modelled as a bin with a capacity equal to the schedule horizon. For casting it is similar with each casting machine a separate bin. Each moulding or casting activity uses an amount of space in the bin according to its duration. In contrast to the classical bin packing, instead of the number of bins we are minimising here the sum of object sizes in a bin.

The difference between the IP models for moulding and casting is that in moulding we have to take into account the time for changeovers between different types of mould being made. Each moulding activity has a type associated with it and we therefore keep track of the number of different types of moulds assigned to each bin. We can then reassemble the activities in a bin ending up with the least number of changeovers, without affecting the overall duration or utilisation of the cavities. The IP model is described as follows.

Let there be n activities, m cavities, and l mould types. Let also

$$x_{ij} = 1, \text{ if activity } i \text{ is assigned to cavity } j$$
$$x_{ij} = 0 \ \text{ otherwise,}$$

where $i \in \{1..n\}$, $j \in \{1..m\}$.

Let

$$y_{jk} = 1, \text{ if cavity } j \text{ produces mould type } k$$
$$y_{jk} = 0 \text{ otherwise,}$$

where $j \in \{1..m\}$, $k \in \{1..l\}$.

Each activity can only be assigned to one cavity; therefore

$$\sum_{j=1}^{m} x_{ij} = 1, \forall i \in \{1..n\}.$$

Consider that if activity i of type k is assigned to cavity j, then this cavity processes this mould type:

$$y_{jk} \geq x_{ij}, \forall j \in \{1..m\}, \forall i \in \{1..n\},$$

where k is mould type of activity i. On the other hand, if cavity i does not process mould type k, this cavity has no activities of type k:

$$y_{jk} \leq \sum_{i=1}^{n} x_{ij}, \forall j \in \{1..m\}, \forall k \in \{1..l\}.$$

We now introduce an integer variable M to represent the makespan such that the sum of activity durations and setup times S on each cavity is less than or equal to it:

$$\sum_{i=1}^{n} x_{ij} \times duration[i] + (\sum_{k=1}^{l} y_{jk} - 1) \times S \leq M, \forall j \in \{1..m\}.$$

5 Process Design Constraint-Based Architecture

To model and analyse the different process design scenarios, we develop a prototype design workbench based on an architecture illustrated in Fig. 2. This architecture is envisaged with many design alternatives.

An implementation of this allows different manufacturing process designs to be evaluated through a flexible user interface on top of constraint-based solvers. The core of the prototype is based around Microsoft Excel that stores the process design components and configurations allowing them to be brought together in appropriate ways. It also presents the results in a manner clear to the user, in this case the manufacturing process planner. Excel as an interface medium has the advantage that it is readily accepted as standard in many industrial organisations. In addition, it has a sufficient level of functionality in statistical analysis and representation of data.

Within the system, the process design is converted into an optimisation data file and passed out to ILOG OPL Studio [6], a constraint-based problem solving technology containing pre-compiled models. The choice between IP and CP-based models is automated depending on the type of scheduling problem presented.

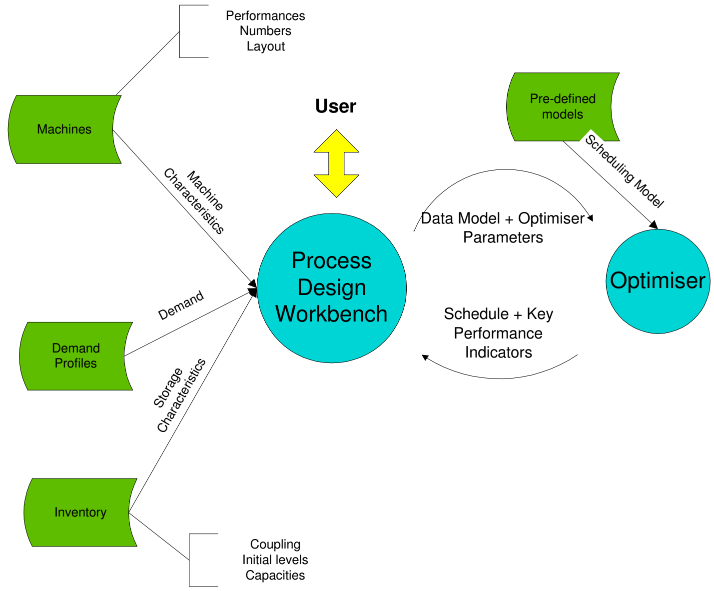

Fig. 2. *An architecture for manufacturing process prototyping.*

6 Experiments

The various coupling scenarios were evaluated on three sets of weekly demand data, which had already been scheduled and produced by the company. The instances differed in the number of orders (2, 60 & 175), but they were approximately the same in overall quantity of production. In all instances, the due date for the orders was 7 days. For the partially coupled scenario, an amount of stock equivalent to 3 days casting was used and for the fully decoupled model an amount equivalent to 7 days was used. 0 days stock corresponded to the fully coupled model. The interface and model allow any number of days stock to be considered.

For comparison between scenarios, we consider that the moulding process is finished when all moulds have been made including the 2 days stabilisation delay. We also assume to be replenishing exactly as many moulds as we remove for casting. In the case of the partially coupled or the decoupled scenario, we also assume that we start the schedule with the moulds ready for casting, so they have been there for a minimum of 2 days.

In practice, the company keeps buffer stock and so moulds are ready to cast at all times. For that reason, the company views the results as excluding the technological delay in moulding for stock (decoupled and partially coupled scenarios), which may result in perceived makespans of up to 2 days shorter. These 'company viewed' results are also reported.

Stock, days	Data Set		
	i	*ii*	*iii*
0	7.27 (**7.27**)	6.84 (**6.84**)	7.43 (**7.43**)
3	7.73 (**5.73**)	7.60 (**5.60**)	7.58 (**5.58**)
7	7.10 (**5.10**)	7.60 (**5.60**)	7.57 (**5.57**)

Table 1. CP model: Makespan, days (with buffer stock).

Stock, days	Data Set					
	i		*ii*		*iii*	
	Moulding	Casting	Moulding	Casting	Moulding	Casting
0	97.5%	84.9%	87.0%	90.5%	90.2%	83.1%
3	86.8%	84.8%	87.0%	79.2%	90.1%	83.4%
7	97.5%	99.3%	87.0%	99.5%	90.2%	90.3%

Table 2. CP model: Average Resource Utilisation.

The results of both models are presented below in terms of the overall makespan, the resource utilisation and the release time of the final finished product. Resource utilisation is measured as percentage of the machine hours used over the total available machine hours within the duration of that process (casting or moulding). The CP model was used in all the coupling scenarios, while as explained before, the IP model was used only for the decoupled processes.

In the experiments, we used ILOG OPL Studio 3.7 which contains both CP and IP solvers. For the IP model, the solver was run with its default settings.

6.1 CP Model Results

In all the cases examined, a first solution was obtained within 5 minutes. The CPU time limit was chosen to be 1800 seconds. The makespan reported covers the overall duration from the start of the first activity to the end of the last activity, which may be either a moulding or a casting activity. The results of the CP model are shown in Tables 1, 2 & 3.

The results first show that orders were satisfied within seven days for all scenarios except the coupled one. In the coupled case, the casting starts too

Stock, days	Data Set		
	i	*ii*	*iii*
0	7.27	6.84	7.43
3	5.28	5.53	5.42
7	4.50	4.40	5.00

Table 3. CP model: Casting Finish Times, days.

late to achieve the order deadlines. The coupled process though has the shortest makespan in two of the three cases. Here there is an efficient use of moulds and a tight interaction between the processes. In the decoupled case, there is a two day delay after the final moulding activity which dictates the overall makespan, although the final product has already been produced. In terms of the buffer results, which correspond more to the actual processes on the shop floor, we see a different pattern with the de-coupled scenario coming out with a smaller makespan. This though requires more stock to achieve.

The de-coupled scenario also provides better utilisation of both moulding and casting resources across all three datasets.

The partially coupled scenario does not provide any compromise in terms of makespan and utilisation. In a sense it has the worse of both worlds, there is a delay in casting caused by waiting for moulds to become available after the stock has been used up and also a delay at the end to achieve stabilisation of moulds.

In Table 2 we see that for the partially coupled scenario, casting resource utilisation degrees are lower than the respective figures for the fully coupled or fully decoupled scenarios. This lower figure is due to a time gap.

Typical casting resource usage profiles are displayed in Figs. 3 to 5 for all the three scenarios. In Fig. 3, casting does not start until the moulds are ready, some time after the start of the overall schedule. In the partially coupled case, Fig. 4, we can see casting both from stock and from the moulding process. In the fully decoupled scenario, Fig. 5, resource utilisation is independent of moulding and can start immediately from stock. For moulding, the profiles of utilisation are more complex (Fig. 6). Among the moulding activities, those for casting are scheduled first followed by those for stock. This strategy results in a short overall makespan, but the utilisation in some cases goes down. We are in effect imposing an ordering of moulding activities on a cavity such that moulding activities for stock are always scheduled after those for casting, which may mean a less than efficient allocation of resources, resulting in some long activities being scheduled late.

A decomposed CP model is also available to use when the decoupled scenario is detected. This gives the opportunity to optimise each process in isolation. The results are reported in Tables 4 and 5.

The results confirm that the moulding process dictates the overall makespan of the combined model. In this decomposed model it is now possible to improve on the duration of casting. This is reflected in the slightly improved utilisation of the casting resources in data set i.

6.2 IP Model Results

In Tables 6 and 7 we show the results of experiments on moulding and casting within a decoupled scenario using an IP model and solver. The CPU time limit was again 1800 seconds for every instance.

The results show a slight improvement over the equivalent CP decomposed model, with reduced makespans and increased resource utilisation. With this

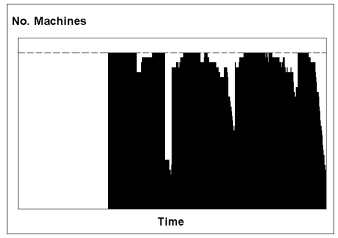

Fig. 3. *Casting Resource Usage Profile, Fully Coupled Process.*

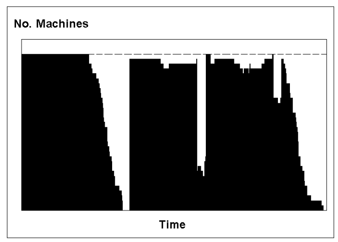

Fig. 4. *Casting Resource Usage Profile, Partially Coupled Process.*

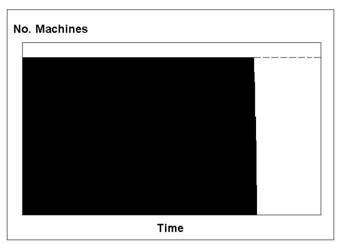

Fig. 5. *Casting Resource Usage Profile, Decoupled Process.*

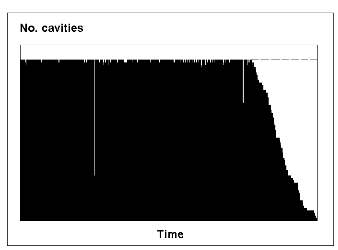

Fig. 6. *Moulding Resource Usage Profile, Partially Coupled Process.*

Data Set	Makespan, days	Average Resource Utilisation, %
i	7.10 (**5.10**)	97.5
ii	7.60 (**5.60**)	86.9
iii	7.57 (**5.57**)	90.2

Table 4. Decomposed CP model, Moulding: Decoupled Makespan and Utilisation.

Data Set	Makespan, days	Average Resource Utilisation, %
i	4.50	99.4
ii	4.40	99.5
iii	5.00	90.3

Table 5. Decomposed CP model, Casting: Decoupled Makespan and Utilisation.

Data Set	Makespan, days	Average Resource Utilisation, %
i	7.08 (**5.08**)	98.1
ii	7.58 (**5.58**)	87.5
iii	7.56 (**5.56**)	90.3

Table 6. IP Model, Moulding: Makespan and Utilisation.

Data Set	Makespan, days	Average Resource Utilisation, %
i	4.48	99.7
ii	4.40	99.6
iii	5.00	90.3

Table 7. IP Model, Casting: Makespan and Utilisation.

set of data the approaches are essentially equivalent in real terms. However, on other data sets it may be useful to keep this alternative model in consideration.

7 Conclusions and Future Work

In this paper, we attempt to develop an understanding of the relationship between manufacturing process design and quality of schedules on a small set of real-world test examples. To do so, we propose a prototype architecture that allows selecting and configuring various types of casting/moulding process. Here we focused on one set of alternative designs, relating to the degree of coupling between the two processes. This approach does not replace detailed scheduling by an experienced scheduler. Instead we aim to produce good schedules in reasonable time. These can then be compared to others in showing the relative merits of different designs.

The results of our empirical study show that for the company, it is possible to move to a lower level of inventory and still meet production demands. However, this re-design comes at the cost of greater perceived risk. More generally, we have verified in one case the usefulness of this architecture for evaluating process designs of moulding and casting.

This architecture is capable of incorporating other design components. The current prototype design workbench already has ideas directly from the company and from the literature. Therefore we are considering additional design components including predefined routes or lines from moulding to casting machines, as this may be dictated by the layout of the shop floor or may be due to compatibility of technical specifications between machines. We also consider variable numbers of machines and performances. Furthermore, we will allow the planner to pre-allocate part of the schedule by associating one order with a set of three machines (2 moulding and 1 casting).

References

1. P. Baptiste, C. Le Pape, W. Nijten. Constraint Based Scheduling. Applying Constraint Programming to Scheduling Problems. *Kluwer Academic Publishers*, 2001.
2. R. Bixby, M. Fenelon, Z. Gu, E. Rothberg, R. Wunderling. MIP: Theory and Practice — Closing the Gap. In M. J. D. Powell and S. Scholtes, editors, System Modelling and Optimization: Methods, Theory, and Applications, pp. 19–49. Kluwer Academic Publishers, 2000.
3. K. Darby-Dowman, J. Little. Properties of Some Combinatorial Optimisation Problems and their Effect on the Performance of Integer Programming and Constraint Logic Programming". INFORMS Journal of Computing, Vol. 10, No 3, pp. 276-286, 1998.
4. M. Garey, D. Johnson. Computers and Intractability; A Guide to the Theory of NP-Completeness, 1979.
5. H. Katayama, D. Bennett. Agility, adaptability and leanness: A comparison of concepts and a study of practice. International Journal of Production Economics, 60-61, 43-51, 1999.
6. I. J. Lustig, J. F. Puget. Program Does Not Equal Program: Constraint Programming and its Relationship to Mathematical Programming. Interfaces Journal, Vol.31 (6), Dec. 2001, pp.29–53.
7. R.N. Nagel, R. Dove, S. Goldman, K. Priess. 21st century manufacturing enterprise strategy: An industry lead view. Iacocca Institute, Lehigh University, Bethlehem PA 1991.
8. N. Sadeh, D. Hildum, T. Laliberty, J. MacA'Nulty, D. Kjenstad, A. Tseng. Blackboard Architecture for Integrating Process Planning and Production Scheduling. Current Engineering: Research and Applications, vol. 6(2), 1998.
9. B. O'Sullivan. Constraint-Aided Conceptual Design Professional Engineering Publishing, December 2001.
10. H.P. Williams. Model Building in Mathematical Programming. *John Wiley and Sons*, 1999.