Journal of Scheduling manuscript No.
(will be inserted by the editor)

On Supporting Lean Methodologies using Constraint-Based Scheduling

Roman van der Krogt[1], John Geraghty[2], Mustafa Ramzi Salman[2], James Little[1]

[1] Cork Constraint Computation Centre
Department of Computer Science
University College Cork, Cork, Ireland
[2] Enterprise Process Research Centre
Mechanical & Manufacturing Engineering
Dublin City University
Dublin, Ireland

The date of receipt and acceptance will be inserted by the editor

Abstract Lean Manufacturing –often simply referred to as "Lean"– is a process management philosophy that aims to improve the way in which products are manufactured. It does this through identifying and removing waste and creating a smooth transition between stages in the production process. To a large extent, it relies on visual and simple mechanical aids to assist in improving manufacturing effectiveness. However, when it comes to combining several aspects of Lean or when dealing with complex environments, quantitative modelling becomes essential to achieve the full benefits of Lean.
In this paper, we show through two detailed case studies how various aspects of Lean can be supported using (constraint-based) scheduling tools. One study concerns a planning support tool to evaluate different Lean initiatives; the other supports the day-to-day scheduling of a complex, Leaned production process

Key words Lean Manufacturing, Constraint Programming, Scheduling

1 Introduction

"Lean" is a process management philosophy that is successfully applied in a variety of sectors worldwide. The term was first used by Krafcik (1988), a principle researcher in the International Motor Vehicle Program (IMVP) at Massachusetts Institute of Technology (MIT), to describe what today is known as the Lean Manufacturing or the Lean Production paradigm (Papadopoulou and Özbayrak, 2005). A major output of the IMVP research efforts was the publication of the book *'The Machine that Changed the World: The story of Lean production'* (Womack et al., 1990), which documents the evolution of the automotive industry from Craft Production, through Mass Production to ultimately Lean Production. Lean is a philosophy that seeks to improve activities by exposing and eliminating all types of waste from the system under consideration, such as a value stream, a manufacturing process, or even a routine job. To do so, it asks an organisation to consider all activities as a series of processes that can be "optimised" and aligned to continuous improvement programmes (Bhasin and Burcher, 2006).

Traditionally, the Lean toolbox has relied on simple mechanical or visual aids to analyse the business processes. However, practitioners have started to realise that quantitative modelling, for example using simulation or mathematical models, can provide a valuable enhancement to the Lean experience and combine several aspects of Lean in conjunction (Detty and Yingling, 2000; Schroer, 2004; Abdulmalek and Rajgopal, 2007). The use of quantitative models in the form of advanced scheduling tools can be used to confirm that Lean motivated changes do bring about the expected rewards, and to support the implementation of such changes.

The purpose of this paper is to illustrate the relationship between Lean Manufacturing and quantitative modelling in the form of scheduling. Our dual goal is to show how Lean practitioners can use scheduling to enhance their work, and to introduce practitioners with a scheduling background to an important possible application area of their work. As Lean Manufacturing matures to more companies and wider problem types, it becomes increasingly important that the available scheduling tools adequately support the most important Lean concepts.

The paper is based around two case studies that illustrate the possible integration of constraint-based scheduling tools within a Lean methodology. The first case study shows how constraint-based scheduling can be used to investigate the impact of proposed Leaning (in particular the reduction of buffers) has on the "schedulability" of a set of processes. The second study demonstrates the support of modelling and scheduling of day-to-day operations within a Cellular Manufacturing environment. These case studies demonstrate how particular Lean concepts can be modelled using constraint-based techniques. The choice of Constraint Programming (CP) was influenced by several factors. Firstly, the problems are scheduling ones in which time needs to be represented at the minute level. Therefore, a CP or Integer Programming (IP) based approach would be considered. Once we analysed the problems more we could see that there was a variety of complex scheduling constraints which would have been too onerous to model in IP; the high level scheduling concepts of CP-based scheduling provided the ideal medium for us.

The remainder of the paper is structured as follows. The next section gives an overview of the Lean Manufacturing philosophy, and is followed, in

Section 3, by a discussion on the special issues that arise when scheduling in a Lean environment. Then, in Sections 4 and 5, the two case studies are presented. The paper ends with a discussion and conclusions on the suitability and benefits of the combined approach of using Lean techniques to analyse a process and suggest improvements with CP-based tools to support those changes.

2 Lean Thinking

The term Lean Manufacturing covers a wide range of techniques.[1] We will discuss some of these techniques in detail in the context of the two case studies. Here, we briefly discuss some of the ideas behind Lean to give the reader an initial appreciation of the approach.

The basic premise of Lean Thinking is that it provides opportunities for organisations to increase productivity, reduce waiting times, lower costs and improve services, if it can be successfully applied to their business sector. The literature provides many examples of successful applications of Lean in a variety of sectors all around the world (Liker, 1998; Siekman, 2000; Zimmer, 2000; Lewis, 2001; Prizinsky, 2001; Strozniak, 2001; Drickhamer, 2002; Bateman, 2002; Rea, 2002; Trombly, 2002; Parker, 2003). One of the trademark features of Lean is its reliance on visual and simple mechanical aids to improve manufacturing effectiveness. Most fundamentals of Lean, such as Value Stream Mapping (Rother and Shook, 1999; Duggan, 2002), Kaizen (Massaki, 1986), and Kanban (Liker, 2004), can be implemented without any major investment in automation.

Womack and Jones (1996) identified the five basic principles underpinning Lean Thinking as:

1. specify value by product;
2. identify the value stream for each product;
3. make value flow without interruptions;
4. pull value from the manufacturer; and
5. pursue perfection.

Lean can therefore be summarised as a methodology for developing a value stream across all products which eliminates waste in the waiting times, transport, inventories, and defects. Lean focuses also on achieving a level production schedule (Krishnamurthy and Yauch, 2007). The techniques that Lean provides to realise these improvements can be split into two categories: technical and cultural (Bhasin and Burcher, 2006). The former category includes tools such as:

Kaizen organising a process of continuous improvement;
Kaikaku a radical overhaul to an activity to improve it (as opposed to Kaizen, which is performed continuously);

[1] In total, Lean offers over one hundred tools (Rowlands et al., 2004).

Cellular Manufacturing a method to arrange factory floor labour into semi-autonomous and multi-skilled teams, or work cells, who manufacture complete products or complex components;

Kanban a signalling system to trigger actions (typically the movement, production or supply of a unit in the factory);

Value Stream Mapping a tool to analyse the flow of materials and information required to bring a product or a service to the customer; and

Single Minute Exchange of Dies (SMED) an approach to drive down setup times.

Companies are encouraged not focus on one or two tools of Lean in isolation. Instead, it is important that companies practise most, if not all, of these improvement techniques, as they focus on different types of waste (Bhasin and Burcher, 2006). In addition to these technical requirements, changes are also required in corporate culture to ensure successful implementation of Lean. These include (and are not limited to) decision making at the lowest level, a strategy of change with clear communication of how the goals will be achieved, clear roles and responsibilities, the development of supplier relationships based on trust and mutual respect, a reduction of schedule changes to maximise stability, and long-term commitment (Bhasin and Burcher, 2006).

A more extended categorisation is given by Shah and Ward (2003), who characterise Lean methods as belonging to one of four categories: namely JIT (just-in-time), Total Productivity Maintenance, Total Quality Management and Human Resource Management. Papadopoulou and Özbayrak (2005) provide a classification scheme for research activity in Lean Production.

It is important to realise that Lean should extend throughout the product life cycle. The paradigm ranges from Lean Product Development, to Lean Procurement, to Lean Manufacturing to Lean Distribution (Karlson and Ahlstrom, 1996). In fact, the application of Lean Thinking outside the boundaries of the organisation is advocated, termed the *Lean Enterprise* by Womack et al. (1990).

3 Scheduling in a Lean Environment

Production scheduling is a critical activity in manufacturing. It concerns the distribution of scarce resources, usually machines, to tasks over time (Pinedo, 2002). The importance of quantitative modelling for scheduling is no less in Lean environments than it is in traditional environments, but the emphasis of the role of scheduling differs. Scheduling in a Lean environment, plays two important roles. Firstly, it has a role in developing improved strategies for dealing with uncertain demand. The aim is to reduce variability in order to dampen the impact of fluctuating demand on the plant, which is costly in terms of holding costs or customer dissatisfaction. The second important role of scheduling in a Lean environment is the levelling of the schedule by smoothing out the peaks and valleys in the production schedule, in order to try to run a fixed quantity of all the operations. Notice the difference between the objectives: the former tries to deal with uncertainty, whereas the second deals with the seasonal ebbs and flows of order volumes. The benefits of a constant production and levelled schedule are reduced overall waste. This reduction can be in the form of fewer operators standing idle while waiting for work, or fewer machines and tools that require high investment sitting unused (McKellen, 2004).

These two aspects, however, cannot be the only goals. One of the problems of focusing on one aspect of production is that there are so many other inter-relationships present. For instance, to achieve a level production with variable demand may mean over-producing some product or shorting some customers. This means that there is always a trade-off that has to be considered. Quantifiable models provide a way to underpin Lean operations and make sure that no opportunities for efficiencies are missed, or worse, that a problem is created in another part of the factory.

Small-scale Lean organisations may use spreadsheets to schedule their production. Such an approach can be helpful in a pilot context or small scale entities. However, several studies have found significant drawbacks with the use of spreadsheets. According to (Miller, 1994), the overall cause of these drawbacks is the fact that spreadsheet models are not detailed enough, ignoring important parts of the problem (in particular variability and stochastic behaviour (Chance et al., 1996)). For example, Hood et al. (1989) present an analysis showing that the key performance indicators presented by the spreadsheet models were off by more than 20%, due to ignoring aspects such as equipment reliability and staffing.

However, IT systems are a crucial addition for most organisations and yield significant benefits. These potential benefits have only exploded with the addition of the Internet (Bruun and Mefford, 2004). Hence, there is a requirement to enable the scheduling method to integrate with an organisation's Enterprise Resources Planning (ERP) or Supply Chain Management (SCM) system. Indeed, some vendors are already moving in that direction (SAP, 2008). In the case studies that we now present, this integration was achieved by the tool being able to parse standard reports that could be extracted from the ERP system.

4 Case Study I: Evaluating Inventory Reductions

The first case study involves a supplier of eye health care products. The particular production site produces a variety of contact lenses. These are produced in a two-step process, that produces batches of a standard size (which is the same across the range of lenses). During the first step, *moulds* are produced. For each contact lens, two moulds are required: one for the front (the *anterior* mould) and one for the back (the *posterior* mould). Around 250 different types of lenses are produced from around 20 different anterior mould types, and 20 posterior ones.[2] In the second step (casting), a pair of moulds is injected with the actual lens material. A complicating factor is that the moulds have a specific time window within which they should be used: a minimum curing time *minDelay* is required to stabilise the moulds, and as the quality of moulds degrades over time, a maximum waiting time *maxDelay* ensures the quality of the lenses. This leads to a high risk of waste if the lenses are not used in time.

One of the ways to achieve a reduction in waste advocated by Lean is to strive for zero inventories. Notice that due to the curing times, this is not fully achievable in this context. Hence, a *reduction* in buffer stock was sought. At the same time, the company wanted to ensure that the system remains *robust*: a problem in the production of moulds should not have an immediate knock-on effect with the casting of the lenses. Experts at the factory could not predict the exact consequences of changes in the buffer strategy, and hence could not adequately assess the trade-off between keeping a minimum inventory and a maximum service level. A quantitative model was needed to support the experts and a constraint-based tool was developed (van der Krogt and Little, 2006).

4.1 The Lean Production Environment

Given a quarterly forecast, weekly production volumes are established for the factory. This allows some smoothing to take place, spreading out the load of heavier weeks. The concept depends on an elementary Lean principle of levelled production called *Heijunka* (Ranky et al., 2003; McKellen, 2004; Gregory, 2004; Papadopoulou and Özbayrak, 2005), which involves the levelling of production by both volume and product mix. (A drawback of Heijunka is that if customers want to put in orders at short notice or increase their demand significantly it can have major disruptions, and this is where product stock comes into the production equation. In this case study, we were given the weekly production volumes, and we did not have to reason about this aspect of the problem.)

The case study includes only the two main processes of moulding and casting with a buffer stock in between. The buffer operates on a *pull principle*, a commonly used technique from Lean Thinking that emphasises only replenishing what is used. Pull replenishment works on the basis of segmenting the Work-In-Process (WIP) or finished goods supermarkets into equal units referred to as Kanbans. Once a certain predetermined quantity of Kanban units is consumed, a signal is generated to indicate the requirement to schedule production or authorise a new order for replenishment. In order to allow for a Kanban system to be productively implemented, a few prerequisites are needed:

1. Demand for the items must be steady and continuous;

[2] Hence, not all combinations of anterior/posterior moulds are valid.

2. Replenishment time must be relatively short in comparison to order lead time; and
3. Raw material needs to be readily available to allow immediate production to commence, once a signal is generated.

If demand for the items is not steady and continuous, inventory models including uncertainty can be used to determine the optimal parameters of the Kanban system in terms of replenishment levels and size of the supermarkets.

The second Lean aspect of the production process in this case study is the *pacemaker* process. The pacemaker is usually a critical operation with limited resources or capacity. Scheduling at this one point will result in a pulling effect on work from upstream processes and flowing product through the subsequent processes to the customer. Here, the casting process pulls moulds from the stock, which are then replenished by the moulding process. The Kanban, then, equals the set of moulds required to produce a certain batch of lenses. Depending on the replenishment level, we can operate in a tightly coupled fashion (a minimal buffer is used, and stocks are replenished as soon as moulds are taken from the stock) or more loosely coupled (the buffer has enough moulds for several batches, and only after a number of batches are taken out, the moulds are replenished). Another point of view would be that the moulding stage either works on a manufacture-to-order basis (in the former case), or on a manufacture-to-stock (in the latter). So in either case, once a certain predetermined quantity of Kanban units is consumed, a signal is generated to indicate the requirement to schedule production or authorise a new order for replenishment.

This integration of both Push (in this case of lenses to the customer) and Pull (of moulds to the casting stage) is proposed in the Lean literature by e.g. Hodgson and Wang (1991a,b) and Cochran and Kim (1998). Notice that they assert the need for quantitative analysis to determine the parameters of such integration. Here, constraint-based scheduling is used to support that analysis.

4.2 The Model

The core of the scheduling model is illustrated in Figure 1. The schedule is derived from the forecasted order volume for the week. This defines how many batches of each type of lens are required. We denote the set of all batches by β. Each batch $b \in \beta$ has an associated size $size(b)$, requires anterior and posterior mould types $ant(b)$ and $post(b)$, and produces lenses of type $lens(b)$. We define three activities for each batch $b \in \beta$: one casting activity $Cast[b]$, and two moulding activities (one for the front, $Ant[b]$, and one for the back, $Post[b]$). The activities are carried out on a set of M

moulding machines, denoted by $mouldMach[1 \ldots M]$, and a set of C casting machines, $castMach[1 \ldots C]$.[3] In the model, this is formulated as follows:

$$\forall_{b\in\beta}\, Ant[b] \text{ requires(1) } mouldMach$$
$$\forall_{b\in\beta}\, Post[b] \text{ requires(1) } mouldMach$$
$$\forall_{b\in\beta}\, Cast[b] \text{ requires(1) } castMach$$

As discussed above, the activities can either be tightly coupled or loosely coupled, depending on the level of buffer stock that is kept. In practise, this will mean that some batches can be cast using moulds from stock (which is later replenished), whereas other have to be cast from moulds which are to be produced specifically for the casting. This can be specified as a percentage that is interpreted as follows. For each type of lens, we determine the total volume for that lens type. Then, we calculate how many of these lenses we can cast from stock, as given by the percentage. This determines which batches can be cast from stock, and which not. For example, at 30%, when we have 5 batches of equal size, we can cast one from stock (i.e. decoupled), one partially from stock and 3 will have to be cast in a coupled fashion. For simplicity, the presented model only takes into account batches that are either fully coupled or fully decoupled; this is trivially extended to hybrid batches by splitting those into two smaller homogeneous ones. For each batch $b \in \beta$, $coupled(b)$ denotes whether a batch is to be cast from stock ($coupled(b) = \text{False}$), or not ($coupled(b) = \text{True}$).

In the former case, we assume that all stock is cured, i.e. can be used from the start of the schedule, and remains in good condition over the schedule horizon. All stock that is replenished, is produced in such a way that at the end of the schedule, all stock is available for immediate casting (i.e. by the schedule horizon, they have been ready for at least *minDelay* time). We can model this as follows:

$$\forall_{b\in\beta}\, \neg coupled(b) \Rightarrow [Ant[b].end + minDelay \leq scheduleHorizon]$$
$$\forall_{b\in\beta}\, \neg coupled(b) \Rightarrow [Post[b].end + minDelay \leq scheduleHorizon]$$

When we are not casting from stock, the three activities for each batch are adequately ordered (i.e. moulding before casting), and the time windows are observed. Taking into account only the anterior moulds, we get:

$$\forall_{b\in\beta}\, coupled(b) \Rightarrow [Ant[b].end + minDelay \leq Cast[b].start]$$
$$\forall_{b\in\beta}\, coupled(b) \Rightarrow [Ant[b].end + maxDelay \geq Cast[b].end]$$

[3] Notice that, while not explicitly mentioned in the model, there are setup times involved for changing from one type of mould to another on the moulding machines, and for switching from one type of lens to another on the casting machines. Moreover, machines may operate at different speeds, and hence the duration of activities depends upon the machine used. These details have been omitted from the presented model for clarity of presentation.

Given λ : set of lenstypes
μ : set of mould types, divided in posterior moulds μ_P
and anterior mould $\mu_A : \mu = \mu_A \cup \mu_P$
β : set of batch ids
$size : \beta \to \mathbb{N}$: number of lenses in a batch
$ant : \beta \to \mu_A$: anterior mould required for a batch
$post : \beta \to \mu_P$: posterior mould required for a batch
$lens : \beta \to \lambda$: type of lens produced by this batch
$coupled : \beta \to \mathbb{B}$: batch is coupled or not

Resources $mouldMach[M]$: moulding machines,
$castMach[C]$: casting machines

Reservoirs $anteriorStock[\mu_A]$: stock of anterior moulds
$posteriorStock[\mu_P]$: stock of posterior moulds
$lensStock[\lambda]$: stock of lenses

Activities $Ant[\beta]$: activities for moulding the anterior moulds
$Post[\beta]$: activities for moulding the posterior moulds
$Cast[\beta]$: activities for casting the lenses

Constraints $\forall_{b \in \beta}\, Ant[b]$ requires(1) $mouldMach$
$\forall_{b \in \beta}\, Post[b]$ requires(1) $mouldMach$
$\forall_{b \in \beta}\, Cast[b]$ requires(1) $castMach$
$\forall_{b \in \beta}\, coupled(b) \Rightarrow [Ant[b].end + minDelay \leq Cast[b].start]$
$\forall_{b \in \beta}\, coupled(b) \Rightarrow [Post[b].end + minDelay \leq Cast[b].start]$
$\forall_{b \in \beta}\, coupled(b) \Rightarrow [Ant[b].end + maxDelay \geq Cast[b].end]$
$\forall_{b \in \beta}\, coupled(b) \Rightarrow [Post[b].end + maxDelay \geq Cast[b].end]$
$\forall_{b \in \beta}\, \neg coupled(b) \Rightarrow [Ant[b].end + minDelay \leq scheduleHorizon]$
$\forall_{b \in \beta}\, \neg coupled(b) \Rightarrow [Post[b].end + minDelay \leq scheduleHorizon]$
$\forall_{b \in \beta}\, Ant[b]$ produces($size(b)$) $anteriorStock[ant(b)]$
$\forall_{b \in \beta}\, Post[b]$ produces($size(b)$) $posteriorStock[post(b)]$
$\forall_{b \in \beta}\, Cast[b]$ consumes($size(b)$) $anteriorStock[ant(b)]$
$\forall_{b \in \beta}\, Cast[b]$ consumes($size(b)$) $posteriorStock[post(b)]$
$\forall_{b \in \beta}\, Cast[b]$ produces($size(b)$) $lensStock[lens(b)]$

Figure 1: The basic lens production model

The mould stock itself is modelled as a set of reservoirs: one for each type of anterior or posterior mould. The initial level of these reservoirs determines whether we are dealing with a tightly or loosely coupled scenario, as discussed above. The moulding activities produce mould stock, whereas the casting activities take it away. Again only taking into account the anterior moulds, this leads to:

$$\forall_{b\in\beta}\ Ant[b]\ \mathsf{produces}(size(b))\ anteriorStock[ant(b)]$$
$$\forall_{b\in\beta}\ Cast[b]\ \mathsf{consumes}(size(b))\ anteriorStock[ant(b)]$$
$$\forall_{b\in\beta}\ Cast[b]\ \mathsf{produces}(size(b))\ lensStock[lens(b)]$$

Recall that the main question we want this model to answer, is whether or not a given level of buffers is robust, i.e. allows a valid schedule to be computed. To be precise: if, given an order profile, we can identify a solution that produces all required lenses for the week, we say that the setup is robust for this profile. To answer this question, an objective function is defined to minimise makespan.

$$\text{minimise} \max_{b\in\beta} Cast[b].end$$

Notice that the there is no constraint stating that the maximum makespan is equal to one week (which is the period for which the order data is supplied). The reason for this is that if the particular level of buffer stock does not allow a schedule within the 7 day period, we can still use the produced schedule to investigate the reason for running over time, rather than just report a failure.

Our model allows us to go beyond evaluating different buffer stock configurations. It allows different manufacturing process designs to be evaluated through providing a framework for several types of what-if scenarios. Examples of these are adding machines to the current shop floor, removing them, or altering their characteristics (such as production rates). It also permits the alteration of a number of parameters within the process itself (e.g. the stabilisation times of the moulds), as well as changes to demand patterns. All these scenarios are possible with the basic model described above. Additional constraints are required to study what-if scenarios related to binding specific mould types or lens types to specific machines, or interlining machines (where a moulding machine feeds into a specific casting machine, rather than any possible machine). All these scenarios were raised by the company in discussions around the initial buffer stock question. Let $onMachine(a,m)$ denote that an activity a is running on machine m. The following constraints then denote that a specific lens type $l \in \lambda$ can only be cast on a specific machine m_l, and that this machine can only cast this type of lens:

$$\forall_{b\in\beta\cdot[lens(b)=l]}\ onMachine(Cast[b], m_l) = 1$$
$$\forall_{b\in\beta\cdot[lens(b)\neq l]}\ onMachine(Cast[b], m_l) = 0$$

Notice that since a casting activity can only run on one machine, the former constraint implies that no other machine can be used for this type of lens. These constraints can be extended straight-forwardly to deal with moulding machines, and to allow multiple lens types on a machine, or multiple machine per lens type. Inlining of machines can be modelled as follows. Suppose moulding machines m_a and m_p are to be in line with casting machine m_c, where m_a supplies the anterior moulds and m_p the posterior ones. Then, the following constraints hold:

$$\forall_{b \in \beta} onMachine(Ant[b], m_a) \Leftrightarrow onMachine(Cast[b], m_c)$$
$$\forall_{b \in \beta} onMachine(Post[b], m_p) \Leftrightarrow onMachine(Cast[b], m_c)$$

4.3 The Search Strategy

To estimate the size of the model, we consider that an average week has between 350 and 400 activities to be scheduled. This amounts to roughly 100,000 variables and 400,000 constraints over these, as reported by Ilog OPL Studio 3.7.1, which we used to develop the models (for the basic model, without inlining or dedication of machines to particular lens or mould types). We used OPL Studio for two reasons: firstly, it allows easy integration with MS Excel, and other standard Office software that the users are comfortable with. Moreover, the Management Information System in use by the company has an interface to MS Excel as well, allowing the input data to be easily fetched from the MIS. Secondly, at the beginning of the project, the user did not have a clear view of what aspects of the problem should be incorporated into the model and which not. OPL Studio can be used to rapidly prototype a model for feedback from the user.

To deal with the computational complexity of a problem of such a size, a heuristic search procedure was used, which has been found in practise to find good solutions quickly. This dynamic strategy is based on the SetTimes heuristic (LePape et al., 1994) and consists of 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 this time. After that, it considers two alternatives: to schedule the activity at time t or to postpone it. 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. One aspect of this search procedure, is that it finds an initial solution quickly, and produces more refined solutions over time, i.e. it is an anytime algorithm (Zilberstein, 1996). This way, the user may explore a number of scenarios quickly, while studying promising scenarios more thoroughly by allocating more time to the optimiser. Initial solutions are typically found within one to two minutes on a 2GHz machine.

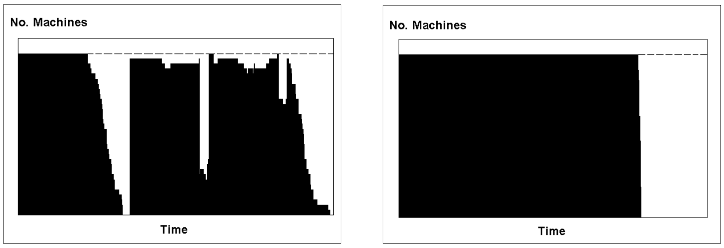

a. Stock shortage b. Sufficient stock

Figure 2: Examples of effects of stock level on casting performance. The vertical axis shows the number of casting machines that are operating at a given time. The dashed line represents the total number of machines. The horizontal axis (time) has the same scale in both graphs.

4.4 The Model in Action

Once the schedule is obtained, the key performance indicators of manufacturing time, utilisation and lens production, are calculated and presented. The tool therefore includes statistical analysis and is able to represent the data as tables, charts and graphs.

It would be possible for the tool to automatically search for the most optimal level of stock. However, this requires that two solutions can be objectively compared to each other, to determine which of the two levels of stock is to be preferred. As there is a large number of performance indicators involved, we could not determine a set of appropriate weights to accurately reflect the company's preferences. Therefore, the tool was designed to support "what if"-scenarios to let the planner make the final judgement.

To determine which level of stock is most appropriate, we thus explored different levels of stock and showed the effect on the resulting schedules. One of the aspects here is the operations of the casting machines. Figure 2 shows the resource profiles for the casting machines for two cases. On the left, we have a minimal amount of stock. As one can see, the casting machines are not fully used throughout, as there regularly are intervals in which not enough moulds are available to keep all casting machines running. On the right-hand side, we have a situation in which enough stock is kept to keep the machines running. Not only does this lead to a much smoother operation of the casting machines, but also the time to produce the required products decreases. By exploring different levels of stock, and weighing the savings gained by a smoother and quicker casting, against the cost of higher stock levels for different scenarios, the company was able to confirm their current stock levels were satisfactory. (The tool showed they could lower stock levels and still achieve an acceptable schedule, but the gains were deemed too small to change from the current levels "that had always worked in the past.")

In our experiments, we used historical data from different time periods, reflecting the differences in order mix and overall volumes. We found that

Table 1: Schedulability for different materials

material	stock level	makespan	moulding util%	casting util%
original	minimal	7.11	96.3	88.2
	50%	5.51	90.5	85.6
	full	5.18	96.3	99.3
new	minimal	-	-	-
	50%	-	-	-
	full	5.68	96.3	99.3

changes in the type of lenses that are produced have a low impact on the schedule quality. This can be explained by the fact that all lenses require two moulding activities and a casting one. While some order mixes may lead to slightly more efficient combinations of moulds that have to be manufactured, the overall impact of this is small. In contrast, the overall order volume does have a large impact. Obviously, high volumes of lenses can only be manufactured when efficient use is made of the resources. Therefore, such high volumes demonstrate potential problems more quickly, whereas the lower ones tend to be more forgiving.

Due to the nature of the tool, additional scenarios can be explored as well. These include the addition or removal of machines to or from the shop floor, as well as evaluating alternative materials for producing the moulds with. For example, when the supplier announced a new type of material to make the moulds with, with different stabilisation parameters (defining the time window within which a mould is to be used) the company could investigate the effect of this new material on the manufacturing process.[4] As we can see from Table 1, it turned out that, perhaps surprisingly, the new material with the "improved" characteristics had a dramatic effect on the schedulability: unless full stocks were kept, there were instances where the new material prevented feasible schedules from being produced.

Here, "full stock" refers to a situation where there is enough stock for all casting activities to take place, "minimal stock" refers to a situation where there is just enough stock at the start of the schedule for the first casting operations to start, but the moulds for the other activities still have to be produced, and "50% stock" means that there is enough stock for half of the casting activities.

The tool has been perceived a valuable aid to the planning department. They have viewed the system positively in addressing the types of problems they are faced with on a regular basis. Previously, they were only able to evaluate a few scenarios when operation conditions changed, due to the manual labour involved, and the problems with understanding the precise consequences. The planning department now have visibility of the schedul-

[4] The difference being that the curing time was significantly decreased, at the expense of having a smaller overall stable window in which the moulds can be used.

ing process and are able to plan the factory operation with some knowledge of the scheduling consequences. The tool has been used to validate historical decisions. In all explored cases, the results of the tool corresponded with the experiences in the factory.

5 Case Study II: Supporting a Cellular Work floor

The second case study involves one of the major players in the field of telecommunications. In particular, it involves one of their sites producing base stations for wireless infrastructure, such as CDMA and UMTS networks. Located between the antennae and the ground network, the function of these base stations is to handle the signals the antennae receive and send. Depending on the model, their size is roughly that of a standard kitchen refrigerator and they are packed with various pieces of electronics: filters, amplifiers, circuit packs, etc. The exact configuration of a cabinet is dependent upon the circumstances it is being placed in: the type of network (e.g. CDMA or UMTS), the frequency (the GSM standard defines eight frequency bands, for example), physical location (inside or outside), network density (is it expected to handle a high or low volume of calls, what is the area the cabinet covers), etc. Some 20 product groups can initially be distinguished. However, as one can understand from the many aspects that are taken into consideration, the variation in cabinets is large, even within one product group. For this reason, they are built to order.

The site that we worked with produced several hundred cabinets per week on average. The production takes place in three stages: assembly, wiring and testing. The durations of each the stages depends on the particular product group of the cabinet.

Assembly The first step in the process starts with a partially pre-populated cabinet supplied by an external manufacturer containing certain basic features such as power supplies, cooling and a back plane. To this, the company adds the required amplifiers, filters, circuit packs and other hardware according to a predetermined schema, according to the customer's chosen configuration.

Wiring The second step involves physically interconnecting the hardware that was added during assembly. Each component has a number of input and output ports that have to be connected with a wire to the outputs and inputs of other modules, again according to a predefined schema.

Testing The final step involves the validation of the completed system. For this, each cabinet is connected to a test station that subjects it to a specially designed set of test signals. If the output of the cabinet is not according to a specification (e.g. due to a cable that is not firmly fixed in place), the test engineer diagnoses the system and corrects the fault.

Of the three stages, assembly is a relatively low skill operation and requires the least amount of training. The wiring step is more complicated, because a great number of connections has to be made between a great number of connectors that all look identical. Moreover, there is a high degree of freedom in the way the connections can be made. Not in terms of the inputs and outputs that have to be connected (these are fixed), but in terms of the order in which the connections are made (i.e. which cables run in front of other cables), the position of the cables (e.g. along the left or the right side of the cabinet) and which cables are tied together. The final step of testing is also complicated, as it involves making diagnoses for the detected anomalies and repairing them. Notice that the high variability in orders exacerbates the complexity in wiring and testing.

5.1 The Lean Production Environment

To address the complexities of the wiring and testing steps, the management decided to introduce a new setup that allows wirers and test engineers to specialise on product groups. The Lean methodology that promotes this organisation of workforce is *Cellular Manufacturing.* Cellular manufacturing, sometimes called cellular or cell production, arranges factory floor labour into semi-autonomous and multi-skilled teams, or work cells, who manufacture complete products or complex components. Properly trained and implemented cells are more flexible and responsive than the traditional mass-production line, and can manage processes, defects, scheduling, equipment maintenance, and other manufacturing issues more efficiently (Irani, 1999). Because the assembly stage did not demonstrate any issues, it was decided that this step would continue to operate as it did before. The resulting layout is depicted in Figure 3. Since the cells are dedicated to certain product groups, the order in which cabinets are assembled becomes much more important. An incorrect sequence may starve certain cells, or lead to building a stock of cabinets waiting to be processed on an already overloaded cell (or indeed both). Throughout the due dates for all the orders must also be taken into consideration. For this reason, a scheduling tool was required (van der Krogt et al., 2007).

Besides the move to a cellular work floor, the site was already implementing various other Lean strategies. Most importantly, goods are produced in a *Just-In-Time* fashion. This applies to the finished products, that, insofar as possible, are produced on their due dates. The principle extends over the entire production line, with each step ideally handing over cabinets as the next station frees up. To this end, the principle of *Takt time* is used to synchronise the stations. Takt time can be defined as the maximum time allowed between two products being produced in order to meet demand. Takt time is estimated prior to generating a schedule. The operations have to be aligned with that rate in order to avoid delays or shortages. However, instances where a production facility is faced with uncertainties such as the arrival of urgent orders, unpredictable machine breakdown or resource shortages, may have an impact on the Takt time. In such cases the Takt time

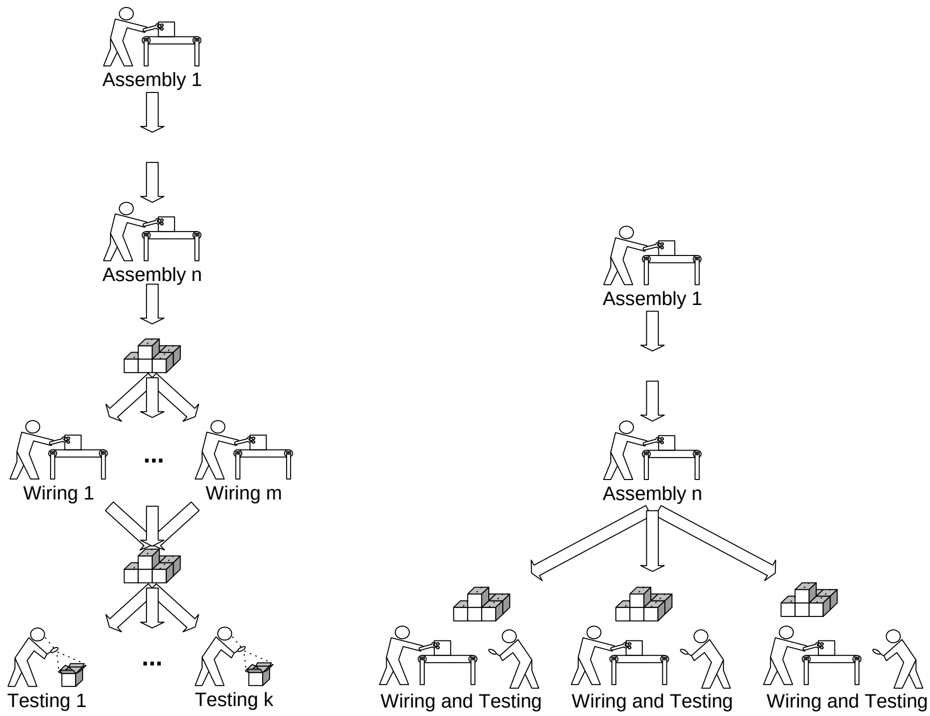

(a) Assembly line (b) Cell-based work floor

Figure 3: From assembly line to cell-based work floor

needs to be recalculated incorporating remedial actions in order to revamp the schedule. This latter point demonstrates again the need of quantitative models to show the implications of changing Takt time in light of uncertainty. First, changing Takt time may mean making changes to resources to adjust their rate of production; is this possible, how quickly can this be done and at what cost? Are the machines involved capable at processing at the new rate? If not, what is the best that can be achieved? Again these are important questions to have answered when adopting a Lean philosophy in a complex environment. Secondly, quantitative models can provide schedules that are robust to change within quantified boundaries (Wu et al., 2006).

Since the production is aligned to a "beat" (i.e. the takt time), the workload for each of the stations the cabinets pass through is roughly equal. Hence, at each beat, cabinets move one step up in the process, and one unit is finished. The standardised durations of wiring and testing are much larger, and hence, one assembly line can feed multiple cells, as we already saw in Figure 3. Notice that the Takt time requires that all stages in the process are aligned to the beat. However, occasionally, the testing of a cabinet may take too long. If this happens, the problematic cabinet is moved to a special testing area that handles these cases. This ensures that even the uncertain testing stage keeps in line with the beat.[5]

There is a natural cycle in the production levels. Towards the end of a quarter the order levels pick up, but then drop sharply at the start of the next. To this end, smoothing is applied by allowing in the model that orders may be produced ahead of their due date if necessary, and some orders may even slip their due dates by a number of days.[6]

5.2 The Model

There are a large number of activities going on in the factory. The assembly line has a number of stages, and each of the cells is also a multi-stage process. However, because of the way in which the processes are organised, we do not have to be concerned with detailed subprocesses. Rather, we can focus on the two main processes, that of assembly and, the wiring and testing (within the cells). The *flow* of cabinets through the assembly line and through the cells means that we can aggregate the stages together into a single activity. Since there is a single assembly line servicing multiple cells, the flow is broken between the two stages, and we have to distinguish between them. We explain how to model this for the assembly line and the cells can be modelled in the same way.

First, observe that the assembly line has n stages. The stages are chosen in such a way, that each requires an amount of work that can be performed within the Takt time. Thus, the assembly line can hold n cabinets at the same time, provided that these are at different stages. We can enforce this property by allowing only one cabinet to enter the assembly line each beat. To this end, we introduce two resources. One unary resource (i.e. one that can be claimed by only a single activity at a time) will be used to space the cabinets in time. A discrete resource with capacity n is used to represent the n stages, as it can be claimed by n activities at any one time. This is illustrated in Figure 4, where we assume $n = 3$.

The upper time line represents the unary beat resource: the first beat, activity 1 may start, then 2, etc. This leads to the staggered use of the discrete assembly resource as shown in the lower resource profile. When the third activity starts, three activities are using the resource. However, they have all started at different beats, and so are at different stages in the assembly line.

[5] Unless, of course, an abnormally high number of faulty units is produced, in which case the fact that the dedicated testing stage has not enough capacity is not the main issue to be resolved. (Instead, attention should be focussed on resolving the cause of these faulty units.)

[6] Buyers are often optimistic when it comes to planning permissions being given to place the cabinets. When planning permissions are not in place yet, they usually do not mind the cabinet being delivered late – it saves them on storage costs.

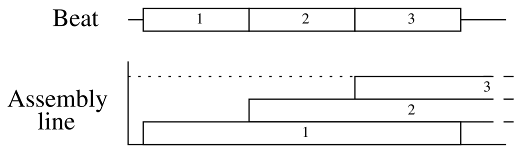

Figure 4: Example of using both a discrete resource and a unary resource to represent a single entity (in this case: the assembly line)

Formally, this is modelled as follows. Given a set π of product types and ω of cabinets to produce, with $o \in \omega$ of type $type(o) \in \pi$, we introduce two activities for each $o \in \omega$: $assembly[o]$ to represent the assembly steps and $assemblyBeat[o]$ to align to the beat. Let $asmblyLine$ be a discrete resource representing the assembly line and $asmblyLineBeat$ a unary resource:

$$\forall_{o\in\omega}\ assemblyBeat[o]\ \textsf{requires}\ asmblyLineBeat$$
$$\forall_{o\in\omega}\ assembly[o]\ \textsf{requires(1)}\ asmblyLine$$

Moreover, we have to align the start times of the two activities to each other and to the takt time τ that determines the rate of production:

$$\forall_{o\in\omega}\ assemblyBeat[o].start = assembly[o].start$$
$$\forall_{o\in\omega}\ assemblyBeat[o].start \bmod \tau = 0$$

As mentioned before, the cells can be modelled in a similar fashion. However, as there are a number of cells, each with its own pair of resources (i.e. one for the beat and one for the cell itself), we have to ensure that we only use these resources in the correct pairs. Let $orderOnCell(o, c)$ denote that a cabinet o is to be scheduled on cell c, then:

$$\forall_{o\in\omega,c\in\gamma}\ orderOnCell(o,c) \Leftrightarrow \big[wireTestBeat[o]\ \textsf{requires}\ cellBeat[c]\big]$$
$$\forall_{o\in\omega,c\in\gamma}\ orderOnCell(o,c) \Leftrightarrow \big[wireTest[o]\ \textsf{requires(1)}\ cell[c]\big]$$

Additionally, we can only assign cabinets to cells that are capable of this product type. Let $cap : \gamma \Rightarrow 2^{\pi}$ be a function returning the capabilities of a cell, i.e. a set of product types that this cell is capable of. Then we have:

$$\forall_{o\in\omega,c\in\gamma}\ orderOnCell(o,c) \Rightarrow type(o) \in cap(c)$$

Finally, each order has a due date. Ideally, all orders will be finished before the due date. The reality of the factory is, however, that due dates will sometimes be missed in periods of high order volume (such as the end of quarters). For each order $o \in \omega$, we introduce a variable $delayed[o]$ that is true iff the due date for o is missed. These variables are used in the objective function, as described below.

$$\forall_{o\in\omega}\ (1 - delayed[o]) \times wireTest[o].end \leq dueDate[o]$$
$$\forall_{o\in\omega}\ assembly[o].end \leq wireTest[o].start$$

Given	π : set of product types
	ω : set of cabinets to produce (the orders)
	$type : \omega \rightarrow \pi$: type of each cabinet
	$release : \omega \rightarrow \mathbb{N}^*$: release time of an order
	$dueDate : \omega \rightarrow \mathbb{N}^*$: due date of an order
	γ : set of available cells
	$cap : \gamma \rightarrow 2^\pi$: capabilities of each cell
	τ : takt time on the assembly line
Variables	$orderOnCell[\omega, \gamma]$: denotes that an order is assigned to a cell
	$delayed[\omega]$: denotes that an order will not make its due date
Resources	$asmblyLineBeat$: unary resource for the assembly line
	$asmblyLine$: discrete resource with capacity A
	$cellBeat[\gamma]$: unary resource for each cell
	$cell[\gamma]$: discrete resource with capacity C
Activities	$assemblyBeat[\kappa]$: activity for synchronising assembly to the beat
	$assembly[\kappa]$: assembly activities
	$wireTestBeat[\kappa]$: activity for synchronising cell activity to the beat
	$wireTest[\kappa]$: cell activities
Constraints	$\forall_{o \in \omega}\ assemblyBeat[o]\ \textsf{requires}\ asmblyLineBeat$
	$\forall_{o \in \omega}\ assembly[o]\ \textsf{requires(1)}\ asmblyLine$
	$\forall_{o \in \omega}\ assemblyBeat[o].start = assembly[o].start$
	$\forall_{o \in \omega, c \in \gamma}\ orderOnCell(o, c) \Leftrightarrow \left[wireTestBeat[o]\ \textsf{requires}\ cellBeat[c]\right]$
	$\forall_{o \in \omega, c \in \gamma}\ orderOnCell(o, c) \Leftrightarrow \left[wireTest[o]\ \textsf{requires(1)}\ cell[c]\right]$
	$\forall_{o \in \omega}\ wireTestBeat[o].start = wireTest[o].start$
	$\forall_{o \in \omega}\ assembly[o].end \leq wireTest[o].start$
	$\forall_{o \in \omega, c \in \gamma}\ onCell(wireTest[o], c) \Rightarrow type(o) \in cap(c)$
	$\forall_{o \in \omega}\ assemblyBeat[o].start \bmod \tau = 0$
	$\forall_{o \in \omega}\ (1 - delayed[o]) \times wireTest[o].end \leq dueDate[o]$
	$\forall_{o \in \omega}\ assembly[o].start \geq release[o]$

Figure 5: The basic cellular manufacturing model

The complete model can be found in Figure 5.

The objective function is a combination of several aspects of the problem. The first and most important aspect is that the number of late orders is minimised. Secondly, we want to minimise the time that cabinets spend waiting in the buffer at a cell. Thirdly, we want to minimise the number of different products that are assigned to a cell while staff gets accustomed to the cellular setup. Finally, the load of the cells should be roughly equal. Let $numProducts[c]$ denote the number of product groups that are scheduled for cell c, $load[c]$ denote the time that c is scheduled to be active, and let $avgLoad = \frac{\sum_{c \in \gamma} load[c]}{|\gamma|}$, then the objective function equals

$$\begin{aligned}
\text{minimise} \quad & f_1 \times \sum_{o \in \omega} delayed[o] + \\
& f_2 \times \max_{o \in \omega} \left[wireTest[o].start - assembly[o].end \right] + \\
& f_3 \times \max_{c \in \gamma} \left[numProducts[c] \right] + \\
& f_4 \times \max_{c \in \gamma} \left[load[c] - avgLoad \right]
\end{aligned}$$

where $f_1, \ldots f_4$ are appropriate weights, of which f_1 carries the most weight, and f_3 reflects the experience of the work force. The actual weights were established experimentally in consultation with the user.

5.3 The Search Procedure

The tool produces weekly schedules. Each week, a few hundred cabinets are produced (on average), resulting in around 150,000 constraints over roughly the same number of variables. Due to the size of problem we have to develop a good heuristic approach.

The cells are the most important aspect to schedule as they are the part which dictates the pace. To reflect this, our search procedure is broken up in two parts. Initially a schedule for the cells is derived, after which the assembly line is scheduled to produce the cabinets for the cells, in a just-in-time fashion. The first search procedure considers each order in turn, ordered by their priorities, the availability of components and their release times. It first tries to add the order to the schedule for the day it should be produced (i.e. its due date). To do so, it tries to assign the order to a cell. It first tries cells that have already one or more cabinets of the same product group assigned. This is to direct the search towards solutions that have a low number of different product groups (and hence changeovers) that are assigned to a cell. If a cell is chosen, the order is assigned the earliest possible start time on that cell during the shift. If a cell cannot be found, the procedure considers producing the order ahead of time, i.e. on a previous day (if possible), or if that is not possible, it accepts that the order runs too late and misses its due date. It is important to note that during this stage we only ensure the existence of a valid schedule for the assembly operations (to ensure that the schedule for the cells is valid). The search procedure does not try to optimise it, as the focus is on maximising the usage of the cells at this stage.

Once the cells have been assigned, a subsequent search procedure is started that optimises the assembly operations. In this second phase, the schedule for the cells is kept fixed, while the maximum time between the completion of assembly of a cabinet and the start of the cell operations on that cabinet is minimised (i.e. the just-in-time aspect is enforced). This greedy algorithm iteratively selects the order for which the minimum possible time that is spent in the buffer between assembly and wiring is currently the largest. For this order, it tries to assign the best possible time left, or removes it from the domain. It does so for all orders, until all orders have been assigned a time to start their assembly.

5.4 The Implementation and Results

The model was implemented in Ilog OPL version 3.7.1 for reasons similar to the previous case study; mainly its rapid prototyping capabilities to prototype a system with the schedulers and production managers, together with its immediate integration through MS Excel to the existing information system. It has been used to schedule the day-to-day operations in the factory until it closed because production was shifted elsewhere.

An important aspect of production monitoring, planning, and control in the real world is the dynamics and the occurrence of exceptions, faults, and new situations. In consultation with the company, we opted for a reactive approach. An important consideration in this decision was that at the time that the project started, the work floor management system did not provide real time information. In short, the basic premise of our approach was that the tool is run at the start of each shift, providing a schedule for that shift. Minor issues are tackled on the work floor, with the scheduling system picking up the changes (if applicable) at the start of the next shift. Major issues (such as a cell suddenly becoming unavailable for some reason) are dealt with by running the scheduling system again.

It is hard to quantify the direct benefit that the scheduling tool has delivered, as it is part of the overall implementation of a different way of working. However, the user observed significant improvements in manufacturing interval, work in process inventory, first test yield, head count, quality and delivery performance. Specifically, the manufacturing interval went down from 16 hrs to 10, and the Work-In-Progress (WIP) inventory decreased by 50%. Although these benefits are mainly achieved because of the Lean motivated change to a cell layout, the scheduling tool is crucial in realising the full potential of this layout. As the user puts it: "*The amount of constraints the program has to handle points out the need for it. Scheduling manually would not allow us to service our customers as well.*"

6 Discussion and Conclusions

Lean Production is a philosophy that has received particularly widespread attention as companies strive for greater efficiencies. In the Lean philosophy, waste (i.e. everything that does not bring value to the customer) is eliminated by striving for zero inventories, zero down times, zero defects, and zero delays in the production process. Lean provides a toolbox of approaches and analysis tools to achieve this, including techniques such as Value Stream Mapping, Kaizen and Kanban. One of the trademark features of Lean is that these are visual or simple mechanical aids and can be implemented without any major investment in automation.

However, as Lean is implemented in more and more complex environments, the traditional tools start to fall short. On the one hand, this is simply a matter of scale. On the other hand, one should also realise that some of the Lean tools may give conflicting recommendations. The more complex a production environment, the harder it is to judge the relative merits of the tools. For this reason, quantitative modelling is becoming more and more important to support and enhance the performance of the tools. Additionally, quantitative models open up new possibilities. In particular, they allow us to accurately predict the consequences of making changes to a process. This is an important ability, as it allows one to weigh the cost of effecting the change with the expected reduction in waste. Moreover, quantitative models may be used to support the day-to-day operations of the new process. This allows a company to fully exploit any new configuration.

In this paper, we illustrate the use of scheduling to support Lean Manufacturing. We see it as providing Lean with a way of enhancing its philosophy to take up the challenge of modern uncertain supply chains instead of trying to control the environment to fit in with known scheduling methods. At the same time, some of the Lean methodologies allow one to take a critical look at the steps in a manufacturing process and their interactions, which may identify an appropriate problem decomposition. In this way, the two approaches complement each other. We presented two case studies that exhibit the strength of a Lean constraint-based scheduling combination. The first study focused on the analysis phase: here, (constraint-based) scheduling helped in understanding the consequences of decisions. The second case saw constraint-based scheduling drive the daily operations of a Leaned work floor. In both cases, the Lean analysis suggested how the model should be developed, through identifying which Lean concepts can be applied. Specifically:

Kanban and Pull Replenishment The first case study shows how Kanban and Pull Replenishment can be implemented using a reservoir and coupled or decoupled activities. The number of coupled / decoupled pairs reflect the tightness of the pull: a tight coupling has a small buffer and requires immediate replenishment, whereas a more loosely coupled scenario allows for a delayed replenishment.

Pacemaker The pacemaker process is the most important process, and the other processes should be organised so as to support pacemaker to work as efficiently as possible. In the two case studies, we implemented the pacemaker by using a custom search procedure that focuses on this process.

Takt time The second case study demonstrated how one can use two resources to force a schedule to synchronise with a beat. For each activity, we introduce a shadow activity that is scheduled on a special unary resource according to that beat. By forcing the two activities to start at the same time, we can align all activities to the Takt time, and ensure that only one activity occurs at each beat.

Flow If one can assume with a high degree of confidence that products will truly flow through the manufacturing process, the second case study shows that it is not necessary to distinguish between the different steps of that process. Instead, the two-resource setup that was used to implement Takt time (using a discrete resource to reflect the number of sequential steps), can be used to model a large part of the production process as one. Only where the flow is interrupted do we need to distinguish between different activities.

Cellular Manufacturing Cells were implemented in the second case study as a regular resource, with constraints specifying which type of products a resource can process. If flow cannot be guaranteed in the cell, a group of resources can be used, where the choice of which cell to use limits which group can be used to perform activities.

Heijunka The second case study implemented production smoothing, by moving orders ahead of time, or allowing orders to slip past their due date, rather than strictly adhering to the required due dates. However, this was only performed over the weekly schedule, not taking into account information that is available for the next week. Truly automated production smoothing should take into account all information that is available, to make a more informed choice about which order to pull ahead and which to let slip.

Just-in-Time The just-in-time aspect of the second case study was achieved by separating the search for a solution into two stages. The first stage of the search produced a feasible solution, which was then optimised for just-in-time aspects in a second iteration.

While this paper serves to demonstrate the strengths of scheduling and, more specifically, Constraint-based reasoning for Lean Manufacturing, we also like to point out areas for improvement. The most important of these is the need for native support of Lean concepts within the constraint-based scheduling framework. While we have shown that it is possible to express Takt time, Pacemakers, Heijunka and Kanban through a combination of regular constraints, special constraints should be designed to be able to model these concepts directly. With these constraints should come powerful propagation methods that allow to efficiently deal with these concepts.

On the Lean side, practitioners have to find novel ways of incorporating the power of (constraint-based) scheduling into their toolbox. "Constrained-based Lean Manufacturing" can have a promising future, but only if people from both communities embrace it.

References

F.A. Abdulmalek and J. Rajgopal. Analyzing the benefits of lean manufacturing and value stream mapping via simulation: A process sector case study. *International Journal of Production Economics*, 107:223–236, 2007.

N. Bateman. Sustainability. In *Lean Enterprise Research Centre Publication*, pages 2–24. Lean Enterpreise Research Centre, Cardiff, UK, April 2002.

S. Bhasin and P. Burcher. Lean viewed as a philosophy. *Journal of Manufacturing Technology Management*, 17:56–72, 2006.

P. Bruun and R.N. Mefford. Lean production and the internet. *International Journal of Production Economics*, 89:247–260, 2004.

F. Chance, J. Robinson, and J.W. Fowler. Supporting manufacturing with simulation: model design, development, and deployment. In *Proceedings of the 1996 Winter Simulation Conference*, pages 114–121, 1996.

J.K. Cochran and S.-S. Kim. Optimum junction point location and inventory levels in serial hybrid push/pull production systems. *International Journal of Production Research*, 36(4):1141–1155, 1998.

R.B. Detty and J.C. Yingling. Quantifying benefits of conversion to lean manufacturing with discrete event simulation: a case study. *International Journal of Production Research*, 38(2):429–445, 2000.

D. Drickhamer. Manufacturer of automotive climate-control systems follows lean manufacturing map. Reprint from Gale Group, 16 October 2002.

K.J. Duggan. *Creating Mixed Model Value Streams: Practical Lean Techniques for Building to Demand*. Productivity Press, New York, NY, 2002.

A. Gregory. Running like clockwork. *Works Management*, 57(2):14–17, 2004.

T.J. Hodgson and D. Wang. Optimal hybrid push/pull control strategies for a parallel multistage system: Part I. *International Journal of Production Research*, 29(6):1279–1287, 1991a.

T.J. Hodgson and D. Wang. Optimal hybrid push/pull control strategies for a parallel multistage system: Part II. *International Journal of Production Research*, 29(7):1453–1460, 1991b.

S.J. Hood, A.E. Amamoto, and A.T. Vandenberg. A modular structure for a highly detailed model of semiconductor manufacturing. In *Proceedings of the 1989 Winter Simulation Conference*, pages 811–817, 1989.

S. Irani. *Handbook of Cellular Manufacturing Systems*. John Wiley & Sons, New York, NY, 1999.

C. Karlson and P. Ahlstrom. Assessing changes towards lean production. *International Journal of Operations and Production Management*, 16 (1):2–11, 1996.

J. Krafcik. Triumph of the lean production system. *Sloan Management Review*, 30(1):41–52, 1988.

R. Krishnamurthy and C.A. Yauch. Leagile manufacturing: A proposed corporate infrastructure. *International Journal of Operations and Production Management*, 27(6), 2007.

C. LePape, P. Couronne, D. Vergamini, and V. Gosselin. Time-versus-capacity compromises in project scheduling. In *Proceedings of the Thirteenth Workshop of the U.K. Planning Special Interest Group*, pages 39–45, 1994.

J. Lewis. Set the stage for success. *Upholstery Design and Management*, 14 (9):1–4, 2001.

J.K. Liker. *Becoming Lean: Inside Stories of US Manufactures*. Productivity Press, Cambridge, MA, 1998.

K. Liker. *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer*. McGraw-Hill, New York, 2004.

I. Massaki. *Kaizen: The key to Japan's Competitive Success*. McGraw-Hill, New York, 1986.

C. McKellen. Production leveling. *Metalworking Production*, November 2004.

D.J. Miller. The role of simulation in semiconductor logistics. In *Proceedings of the 1994 Winter Simulation Conference*, pages 885–891, 1994.

T.C. Papadopoulou and M. Özbayrak. Leanness: Experiences from the journey to date. *Journal of Manufacturing Technology Management*, 16(7), 2005.

V. Parker. Burt's bees implementation of production processes. *Tribune Business News*, (1-3):2–4, 2003.

M. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Prentice-Hall, New Jersey, 2002.

D. Prizinsky. Lincoln looks leaner in its manufacturing process. *Crains Cleveland Business*, pages 1–8, April 2001.

P.G. Ranky, M. Subramanyam, R.J. Caudill, K. Limaye, and N. Alli. Dynamic scheduling and line balancing methods, and software tools for lean and reconfigurable disassembly cells and lines. In *Proceedings of the IEEE International Symposium on Electronics and the Environment*, pages 234–239, 2003.

D. Rea. Following the lean route. Reprint from Gale Group, 20 August 2002.

M. Rother and J. Shook. *Learning to See: Value Stream Mapping to Add Value and Eliminate Muda*. The Lean Enterprise Institute, Brookline, 1999.

D. Rowlands, J. Maxey, M. George, and M. Price. *The Lean Six Sigma Pocket Toolbook: A Quick Reference Guide to 100 Tools for Improving Quality and Speed*. McGraw-Hill, New York, 2004.

SAP. SAP lean planning and operations. `http://www.sap.com/solutions/manufacturing/lean-planning-and-operations-software/index.epx`, retrieved June 23 2008.

B.J. Schroer. Simulation as a tool in understanding the concepts of lean manufacturing. *Simulation*, 80(3):171–175, 2004.

R. Shah and P. T. Ward. Lean manufacturing: Context, practice bundles, and performance. *Journal of Operations Management*, 21(2):129–149, 2003.

P. Siekman. Cessna tackles lean manufacturing. *Fortune*, 141:222–231, 2000.

P. Strozniak. Toyota alters face of production. *Industry Week*, 250(11):46, 2001.

R. Trombly. Running lean running strong: Lean manufacturing processes lead to a stronger, more efficient business. Reprint from Gale Group, August 2002.

R. van der Krogt and J. Little. The PDES workbench. In *Proceedings of the Thirteenth International Conference on Concurrent Engineering: Research and Applications*, pages 619–626, 2006.

R. van der Krogt, J. Little, K. Pulliam, S. Hanhilammi, and Y. Jin. Scheduling for cellular manufacturing. In *Proceedings of the Thirteenth International Conference on Principles and Practice of Constraint Programming (CP-07)*, pages 105–117, 2007.

J.P. Womack, D.T. Jones, and D. Roos. *The Machine that Changed the World: The Story of Lean Production.* Rawson Associates, New York, 1990.

P.J. Womack and D.T. Jones. *Lean Thinking: Banish Waste and Create Wealth in Your Corporation.* Simon & Schuster, New York, 1996.

C.W. Wu, K.N. Brown, and J.C. Beck. Scheduling with uncertain durations: Generating β-robust schedules using constraint programming. In *Proceedings of the ICAPS 2006 Workshop on Constraint Satisfaction Techniques for Planning and Scheduling Problems*, pages 39–45, 2006.

S. Zilberstein. Using anytime algorithms in intelligent systems. *AI Magazine*, 17(3):73–83, 1996.

L. Zimmer. Get lean to boost profits. *Forming and Fabricating*, February 2000.