Workload control under diagnosis

Soepenberg, Gerrit Dinant

Publication date:
2010

Chapter 4.

Workload control dynamics in practice

4.1. Introduction

Make-to-order (MTO) companies increasingly have to cope with fierce competition in today's turbulent markets. Improving and controlling logistical performance is crucial for such companies. Appropriate production planning and control (PPC) can contribute to achieving this aim. Workload control (WLC) is viewed as a production planning and control (PPC) solution specifically designed for the MTO industry (Stevenson et al., 2005). WLC aims to control throughput times by incorporating a restricted release of customer orders to the shop floor, while maintaining an order pool prior to release to buffer against the many uncertainties involved with MTO companies. WLC has been studied extensively in recent decades. It started with a handful of implementations in the eighties/early nineties (Bechte, 1988; Wiendahl et al., 1992). After that several theoretical studies (for example, Wiendahl, 1995; Kingsman, 2000; Fowler et al., 2002; Breithaupt et al., 2002), simulation studies (for example, Melnyk and Ragatz, 1989; Hendry and Wong, 1994; Perona and Portioli, 1998; Bertrand and van Ooijen, 2002; Cigolini and Portioli-Staudacher, 2002), and literature reviews (for example, Wisner, 1995; Bergamaschi et al., 1997; Sabuncuoglu and Karapinar, 1999) have been carried out. However, contemporary empirical studies are relatively scarce and many authors regarded it as a next step (Gaalman and Perona, 2002; Stevenson et al., 2005). Only recently some valuable attempts have been made to foster the implementation of WLC in a practical setting and to the refinement of the concept (see, for example, Stevenson, 2006; Hendry et al., 2008; Stevenson and Silva, 2008).

During the implementations considerable improvements have been observed in practice after WLC has been implemented (for example, Bechte, 1988; Wiendahl et al., 1992). However, performance in simulation studies differs from that seen in reality (Fredendall et al., 2009), a divergence often labelled as the WLC paradox. Nevertheless, the main insights into the functioning of WLC are based on simulation studies, which analyse WLC performance in stationary situations. As a consequence, insights into the functioning of WLC in a dynamic setting, as is generally encountered in practice, are limited.

The main aim of this paper is to contribute to WLC theory by identifying and classifying those issues that arise when WLC is used for controlling performance in a practical dynamic setting. For this purpose, a detailed study in an MTO company, in which WLC is starting to be used, has been carried out. In order to capture the effects of WLC in this practical dynamic setting, detailed quantitative measurements on order progress have been collected and analysed over a period of approximately one year, supplemented with qualitative data obtained in workshops.

The paper starts by discussing the functioning of WLC from a theoretical perspective. Next, the research design and the main analysis techniques are presented. Subsequently, the results section provides an in-depth analysis of the effects of using WLC in a practical dynamic setting, followed by a discussion on the implications of the obtained empirical results for WLC theory. At the end of the paper, the conclusions and opportunities for further research are discussed.

4.2. Functioning of WLC: theory

WLC aims to control the logistical performance of companies. It is regarded as the most appropriate PPC concept for MTO companies (Stevenson et al., 2005). WLC is based on the philosophy that controlling the logistical performance of MTO companies requires a controlled situation on the shop floor, i.e. the throughput times of orders on the shop floor are controlled. Figure 1 visualises the WLC core mechanism, which uses the release decision to control the shop floor throughput times. The core mechanism incorporates the relationships among five distinguishing elements of WLC, as first discussed by Henrich et al. (2004), and their intended influence on the control of shop floor order throughput times.

The first two elements in the core mechanism indicate which point in the order flow and what information are used in performing the third element, central load balancing.

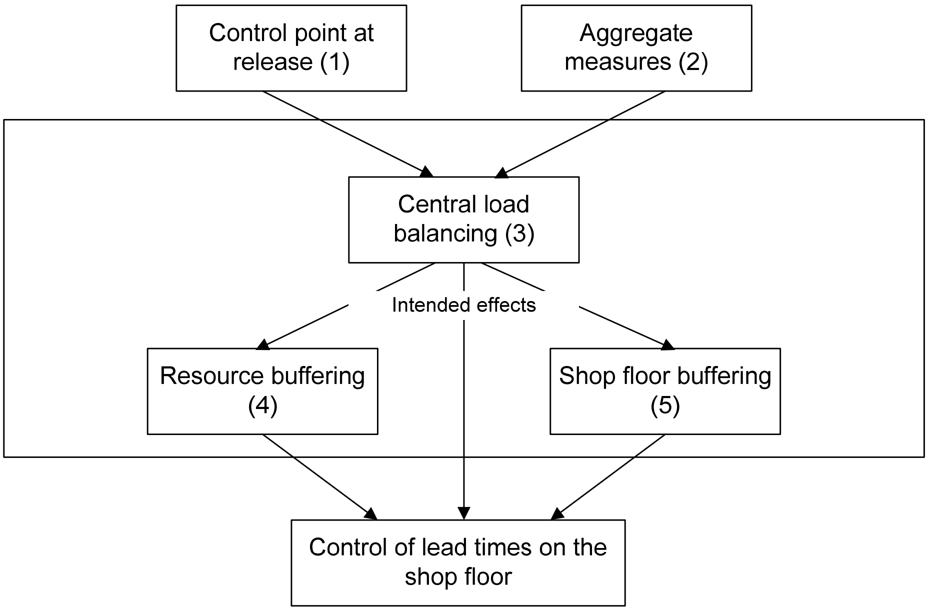

Figure 4.1. Overview of the WLC core mechanism.

(1) Control point is at release
The main decision point for controlling shop floor throughput times is order release.

(2) Aggregate measures
Release decisions are based on aggregate measures, instead of on the detailed information often used in scheduling systems. The use of these aggregate measures avoids nervousness in decision-making in highly variable production environments. The total workload, or a related quantity, released to a resource is commonly used as an aggregate measure for this purpose (see, for example, Bergamaschi et al., 1997).

(3) Central load balancing
Central load balancing refers to the key approach being followed in making release decisions. In balancing loads on the shop floor, these decisions aim to achieve constant buffers in front of the resources on the shop floor. Stable buffers on the shop floor result in controlled and thus predictable shop floor throughput times. In this respect, applying central load balancing enables the determination of appropriate release dates. As such, balancing is regarded important for effective release decisions (see, for example, Germs and Riezebos, 2010). In addition, central load balancing should also take the needs of individual orders into account. The most urgent orders are first considered for release,

which enables a smooth and constant progress of orders on the shop floor and reduces the need for priority changes. The latter relationship is incorporated in the overview shown in Figure 4.1 by means of the direct arrow between central load balancing and the control of throughput times on the shop floor.

The last two elements relate to the intended outcomes of central load balancing: limited and stable buffers on the shop floor (4) and a large buffer in front of the shop floor to absorb fluctuations (5).

(4) Resource buffering
The WIP on the shop floor should be limited. However, because variations in both order processing times and order arrivals occur, it is assumed that some buffering through maintaining controlled small queues of orders is required at the individual capacity groups (Breithaupt et al., 2002).

(5) Shop floor buffering
The largest buffer of orders, the so-called order pool, is placed in front of the shop floor as a whole. The order pool should absorb fluctuations in the flow of arriving orders in order to allow the buffers at the shop floor capacity groups to be small and stable.

The core mechanism can be further supported by output control and order acceptance decisions. For a formal description of all the input and output control decisions the reader is referred to Kingsman (2000). The mechanism outlined in this section provides the core theory on how WLC influences performance and also plays a central role in the research design.

4.3. Research design

The main research question (RQ) of this paper is:

How does the WLC core mechanism function in a practical dynamic setting?

This section discusses the research design established in order to answer this question. The first part of this section discusses the rationale for using a single case study, together with a concise description of the main criteria for and characteristics of the selected company. The second part elaborates on our research approach in this company and the applied research steps.

4.3.1. Case company: rationale and description

This research focuses on a single company which, according to Yin (1989), can be appropriate under several circumstances. The main rationale here is that, in Yin's (1989) terms, it can be regarded as a revelatory case: it provides *"an opportunity to observe and analyse a phenomenon previously inaccessible to scientific investigation."* Here, it relates to a first and unique possibility to gain in-depth empirical insights into how the core mechanism of WLC functions in a practical dynamic setting. In addition, the functioning is analysed over a considerable period of time, which enables insights to be obtained into changes over time; and this longitudinality is another major rationale for a single case approach (Yin, 1989).

The main criterion in selecting a company for this study was that it could be regarded as representative (Yin, 1989) in terms of the application of WLC for controlling logistical performance. As such, an MTO company in a turbulent market was chosen because such a company has to deal with a lot of variety regarding orders. The selected company is a manufacturer of PVC window and door frames and delivers these frames to large building projects and to individual customers. Demand in its market is unstable and subject to all kinds of external influences. For example, progress on large building projects can be affected by both economic and environmental factors. Further, the variety in produced frames is large, with a wide range of sizes, styles and shapes. The production structure in the company can be labelled as a multi-product production line. Although a rather dominant flow does exist within this company, the variety of products results in variable processing times and some routing variety. The routing variety depends on whether doors and opening window elements or other (special) parts as aluminium sills are required in addition to the door frames.

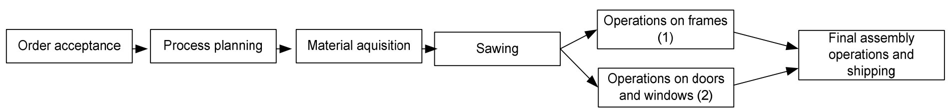

Figure 4.2. Main order flow in company.

The company's general order flow is visualised in Figure 4.2. The first stage in the order flow is the order acceptance process in which new orders are accepted and provisional delivery dates are agreed with the customers. Often, long-term contracts are arranged with companies for orders related to building projects. The second stage is process planning, in which work instructions are created, drawings are prepared and material requirements determined. At the end of the process planning, materials are ordered from suppliers. In this paper we consider the completion of process planning to be the time that an order enters 'the shop floor buffer' (see Figure 4.1), because at that time the required delivery dates are certain. Before process planning, the required delivery dates can easily change because the progress of building projects is hard to predict. After process planning the orders are ready for release to the production department. In the production process the order has to pass through several steps, with resource buffers in front of them.

The starting point in the production process is the cutting to length of PVC bars for the doors or windows. To strengthen the PVC bars, steel parts are sawn to length and fitted into the bars. After this initial sawing process, two parallel lines, one for frames (1) and one for doors/windows (2), can be distinguished, both containing multiple stations. The two lines later converge, after which both doors and opening window elements are installed into the fixed frames and glazing beads are attached to the frames. Special operations, for example the sawing to length of aluminium sills, can be required at various steps of the production line. These special operations are performed in a separate production hall. Finally, orders are shipped to the customer. A detailed overview of all the operations is provided in the Appendix.

4.3.2. Research approach and steps

To gain insights into how the WLC mechanism functions, we carried out longitudinal empirical research in a case company. A considerable period of measurements was undertaken both before (baseline measurement) and during the use (post-change measurement) of the WLC mechanism. The implementation of the WLC mechanism involved cooperation between researchers and company managers and, in that sense, the research can be typified as action research (Coughlan and Coghlan, 2002). The advantage of action research, as a variant of case research, is that acting as a change agent provides deeper insights than can be gained by a case researcher who only records the observed situation (Westbrook, 1995). Although the changes are initiated collaboratively, attempts are made to minimise the influence of the researchers during the baseline and post-change measurements.

Summarising, the research involves three steps (see Figure 4.3). First, baseline measurements (1) are made. This step enables a good understanding of the current way of working. In the baseline measurements, insights are gained into the effect of the currently applied PPC decisions on logistical performance. Based on these results, the WLC mechanism is incorporated during the second step (2). After the changes are established, a set of post-change measurements are collected (3), in which the results of using the mechanism are compared with the earlier results. The remainder of this subsection elaborates on the three research steps.

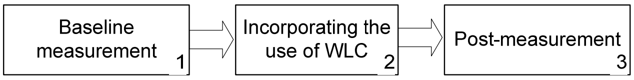

Figure 4.3. Research steps.

1) Baseline measurement

The diagnosis in this study is largely based on quantitative data gathered in the company, supplemented with some qualitative data. In terms of quantitative data, order progress was measured in each relevant step of the realisation process, from acceptance through to delivery. For this purpose, a registration system, developed by a third party, was installed in the company. Using bar code scanning, the order start times and order completion times at each station were recorded. In addition, data on promised delivery times were collected.

Two start-up meetings were arranged once the registration system was installed. In the first meeting, the purpose of the project was explained to members of the managerial staff, including the general manager, the production manager and the production planner. In the second meeting, the shop floor personnel were provided with the more technical instructions on use of the bar code scanners, preceded by a short introduction on the purpose of the project. It was stressed during these meetings that the baseline measurements were intended to reveal the effects on performance of how PPC decisions were currently executed. It was agreed with the personnel that no changes in the PPC would be initiated during this period. Nevertheless, despite all precautions, introducing a registration system could influence the personnel's performance. If personnel know that their performance is being recorded, it can create a Hawthorne effect. To cope with this, and potential other disturbances, two precautions were taken. First, a one month start-up period was included during which the personnel could get used to the system before the data would form part of our study. We were finally interested in changes between the baseline and post-change periods and did not expect the introduction of the scanning system to influence any differences between these periods. Second, the scanning data would also be used by the personnel department to determine working hours and hence salary payments. The thought was that this would motivate the shop floor personnel to ensure the data remained reliable over time.

In the start-up period, the researchers visited the company regularly to help to rectify some start-up problems. Technical assistance was provided by the third party who had developed the registration system. Contact between researchers and managerial staff during the baseline measurement period was limited. Mainly technical assistance over using the scanning system was offered, and no feedback on the results was provided to the company personnel before the end of this stage.

A complete diagnosis based on the quantitative data was carried out by the researchers at the end of the baseline measurement period. The outcomes were discussed with members of the managerial staff in a workshop. During the workshop, additional qualitative data were gathered and analysed to verify the results. Combining the quantitative scanning data with the qualitative data helped to form a rather rich picture of how PPC decisions were currently made. In addition, using multiple data sources enables one to be more convinced over the obtained results (Yin, 1989). Furthermore, this exchange of data between researchers and managerial staff helped to foster a common understanding between researchers and managerial staff of the influence of PPC decisions on the current performance.

2) Incorporating the use of WLC

The aim of the second research step was to adapt the PPC decisions and to find suitable ways to use WLC. The actions performed in this step were taken collaboratively, an important characteristic of action research (Coughlan and Coghlan, 2002). The starting point of the second step was the workshop at the end of the baseline measurement period. Before discussing the results of the baseline measurements, the purpose of WLC and its core mechanism were extensively explained and discussed in the workshop. This enabled a common understanding to be reached of differences between the currently applied PPC decisions and PPC decisions based on the WLC mechanism. After the workshop, several follow-up meetings were organised with the managerial staff. In these meetings, specific directions for improvement, based on the current situation, were elaborated, and the pros and cons of several options collaboratively evaluated. Finally, it was agreed which specific measures would be taken for the next measurement period.

3) Post-change measurement

The final step comprises post-change measurements. For this purpose, data were gathered from the end of the baseline measurement period until the end of the post-change measurement period. As data were recorded continuously, good insights were gained into the transient period immediately after the adaptations. At the end of this post-change measurement period, a new analysis was carried out by the researchers, focused on determining the influence of WLC on the logistical performance of the company. For this purpose, the influence of PPC decisions incorporating the WLC core mechanism were compared with the influence of the PPC decisions in the baseline measurement period. The results obtained were presented in a workshop. Again this workshop aimed to gather additional qualitative data from the managerial staff of the company to improve our understanding of the measured results.

4.4. Applied analysis techniques

Analysing the effects of the currently applied PPC on the logistical performance of the company requires a link to be made between performance and PPC decisions. Amongst others, logistical performance relates to high delivery reliability (see, for example, Wiendahl, 1995). In WLC the key to realising this performance objective is sought in controlling shop floor throughput times (see Figure 4.1). In order to achieve such control, both the average throughput time and its variance need to be monitored. Correspondingly, two types of diagram are used to support our analysis: throughput diagrams and order progress diagrams. The basics of both these analysis techniques will be briefly explained below. It is shown how the use of the core mechanism elements of central load balancing, shop floor buffering and resource buffering are reflected in these diagrams. Further, the support provided by the diagrams in linking these elements to the control of shop floor throughput times is discussed. For a more detailed explanation on throughput diagrams and order progress diagrams, see Wiendahl (1995) and Soepenberg et al. (2008) respectively.

4.4.1. Throughput diagrams

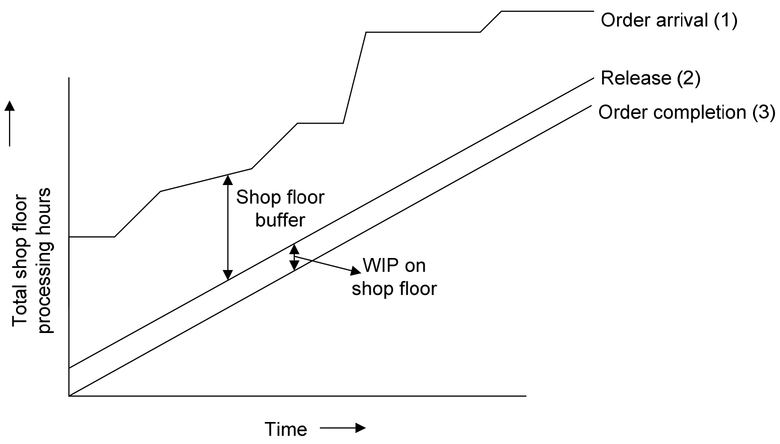

Figure 4.4. Throughput diagram based on WLC core mechanism.

Throughput diagrams can help to analyse the effects of applying the WLC mechanism on average shop floor throughput times. Firstly, adopting central load balancing aims to achieve small and constant buffers on the shop floor. Maintaining these small and constant buffers on the shop floor, i.e. core element 3, should result in uniform but limited WIP on the shop floor. The WIP on the shop floor amounts to the vertical distance between the order release curve (2) and the order completion curve (3) in the throughput diagram. The WIP divided by the completion rate gives an approximation for the average throughput times, which can thus be estimated from the horizontal distance between the curves in the throughput diagram. This relationship is analogous to Little's result (Little, 1961). As such, a steady amount of WIP on the shop floor combined with a uniform order completion rate leads to well controlled average throughput times on the shop floor. In that sense, throughput diagrams are also helpful in analysing how average throughput times change over time.

Secondly, a buffer in front of the shop floor, core element 5, should absorb fluctuations in order arrivals. Figure 4.4 shows how a throughput diagram would ideally look like if this core WLC element is applied. The shop floor buffer is indicated by the vertical distance between the order arrival curve (1) and the order release curve (2). The figure shows how fluctuating order arrivals are absorbed in the shop floor buffer to produce a steady order release rate.

Thirdly, besides maintaining small and stable buffers over time, central load balancing should consider the urgency of individual orders at the order release point. An analysis of the progress of individual orders, and thus consideration of relative urgency, can be performed with the help of order progress diagrams.

4.4.2. Order Progress Diagrams

Order progress diagrams can help to assess the progress of individual orders. The basics of order progress diagrams can be explained with the help of Figure 4.5. This figure shows three curves, representing the progress of three orders.

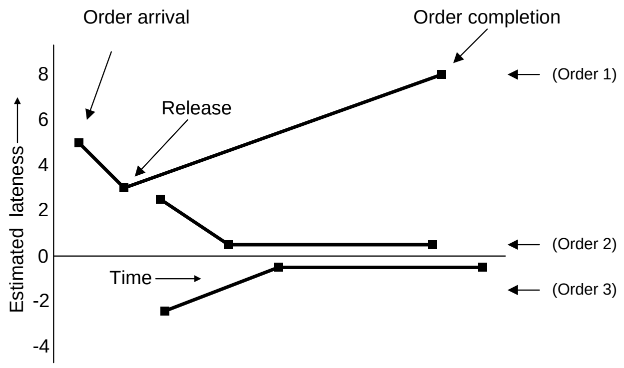

Figure 4.5. Order progress diagram for three orders; Orders 2 and 3 progressed according to WLC.

The horizontal axis in the order progress diagram represents time and the vertical axis indicates estimated lateness. The horizontal position of each dot on a curve represents the time at which a certain stage is completed. The vertical position of the dots on each curve represents the estimated lateness of the respective order on the completion of that stage. The estimated lateness at the end of a certain stage is the lateness an order will have if the throughput times of later stages are equal to the average throughput times for those stages.

Orders with a positive and negative estimated lateness at the order arrival point are commonly referred to as urgent and non urgent orders respectively. Curves with a positive gradient at a certain stage indicate a delay in that stage, while negative gradients indicate orders that are gaining. For example, Order 1 in Figure 4.5 was released relatively early but ultimately delayed due to a relatively long throughput time on the shop floor.

The right end dot of each curve represents the realised lateness of the order on its completion. As such the vertical spread of the curves' right-hand end-points for all orders completed in a certain period gives an impression of the variance of lateness in that period.

Order progress diagrams are particularly helpful in gaining insights into the effects of central load balancing on the progress of individual orders. Central load balancing should lead to predictable shop floor throughput times, which should normally result in a low variance. A low variance in shop floor throughput times would be reflected by flat, horizontal shop floor related segments in the order progress diagram, exemplified here by Orders 2 and 3. This is enabled by having constant and limited WIP on the shop floor, as fewer opportunities then arise for priority changes. Furthermore, releasing the most urgent orders first reduces the need to set priorities on the shop floor. If urgency is taken into account at release, orders that are urgent at acceptance (Orders 1 and 2 in this example) should have short order pool times between order arrival and release, whereas non-urgent orders (Order 3) should be held back.

Combining urgency and balancing considerations should ideally result in converging line segments between order arrival and order release in the order progress diagram. However, achieving an estimated lateness for all orders of less than zero at release is difficult because the WIP level on the shop floor could restrict the release of some orders at their planned release times. This is illustrated by Order 2, where the acceleration through the order pool is insufficient to deliver a negative estimated lateness on its release.

Finally, it should be noted that the order progress diagram shown in Figure 4.5 contains just three orders. For situations in which dozens of orders are in progress at the same time, the visibility of orders in an order progress diagram can be limited. Visibility can be enhanced by showing only relevant subsets of orders in a single diagram.

4.5. Results

This section discusses the results of the empirical research. The two diagrams introduced in the previous section will be used to aid the discussion of results in both the baseline measurement and the post-change measurement periods.

Figure 4.6 shows a throughput diagram for the full research period. Four sub-periods are distinguished, one a baseline measurement (BM) and three periods of post-change measurements (PM). As in Figure 4.4, three curves are drawn: the order arrival curve (1), the order release curve (2) and the order completion curve (3). At the top of the figure, the average weekly input/output rates related to these curves are specified for each period, all in hours per working week. The WIP in the order pool, i.e. the vertical distance between curves 1 and 2, and the WIP on the shop floor, i.e. the vertical distance between curves 2 and 3, are visualised in the lower segment of the figure. Above this component, statistics on the average levels of WIP in the order pool and on the shop floor during each period are provided, all in processing hours.

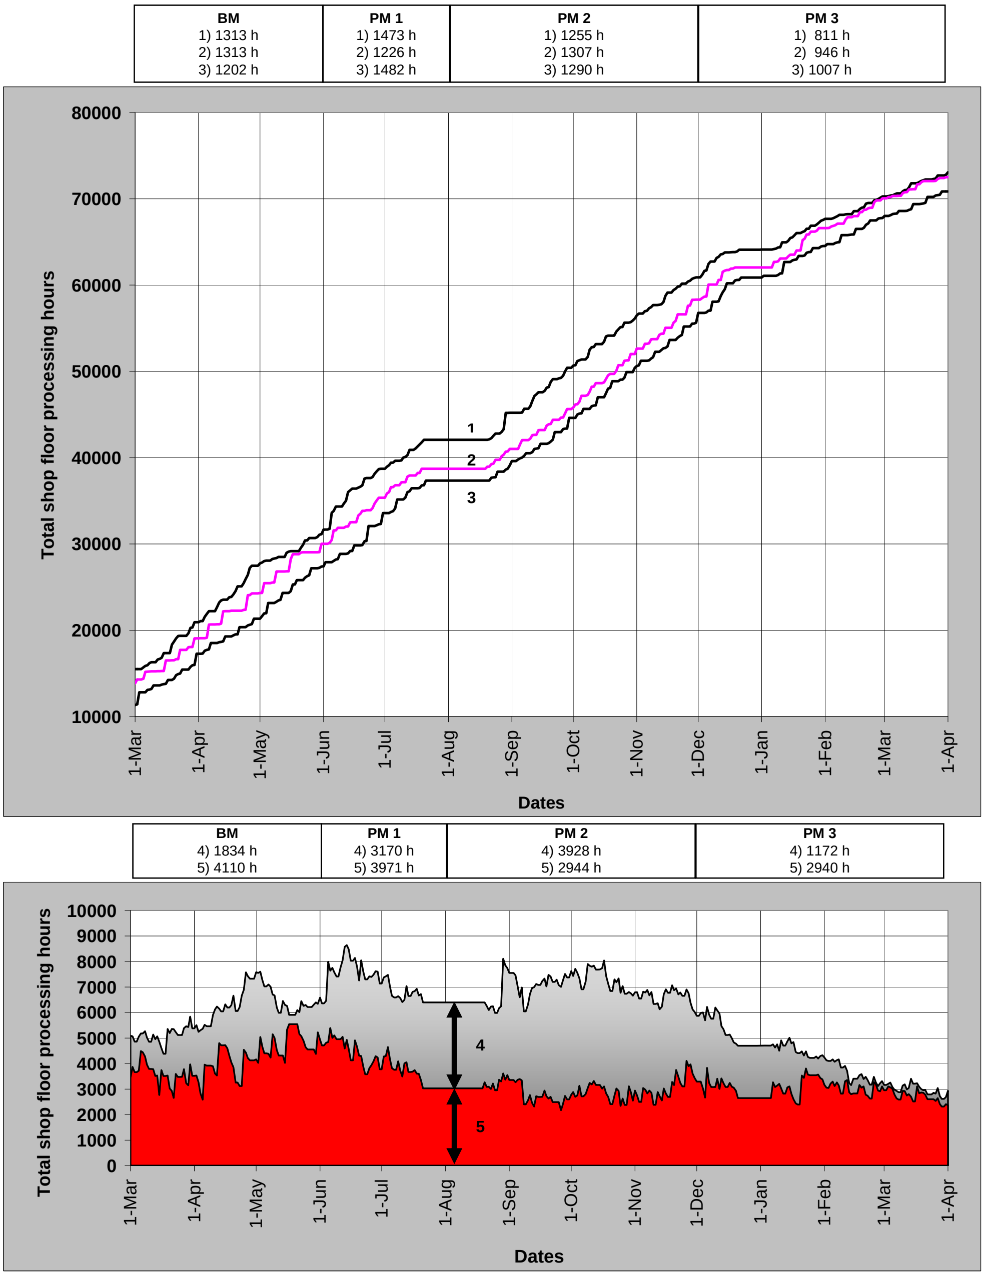

Figure 4.6. Throughput diagram of whole research period. 1) Order arrival; 2) Order release; 3) Order completion; 4) WIP in order pool; 5) WIP on shop floor. For curves 1 to 3, the average weekly input/output level (in hours) is specified per period. On top of the WIP graph, the average WIP levels (in hours) are specified per period.

Figure 4.7 shows order progress diagrams for each of the above periods. The four figures on the left show the progress of the individual orders in the order pool. The starting point of each curve on the horizontal axis is the order arrival date and the end point is the order release date. The three figures on the right show the progress of the same orders on the shop floor. The starting point of each curve on the horizontal axis here is the order release date and the end point is the order completion date. It should be noted that, in order to improve visibility, only orders released in even numbered weeks are depicted in Figure 4.7.

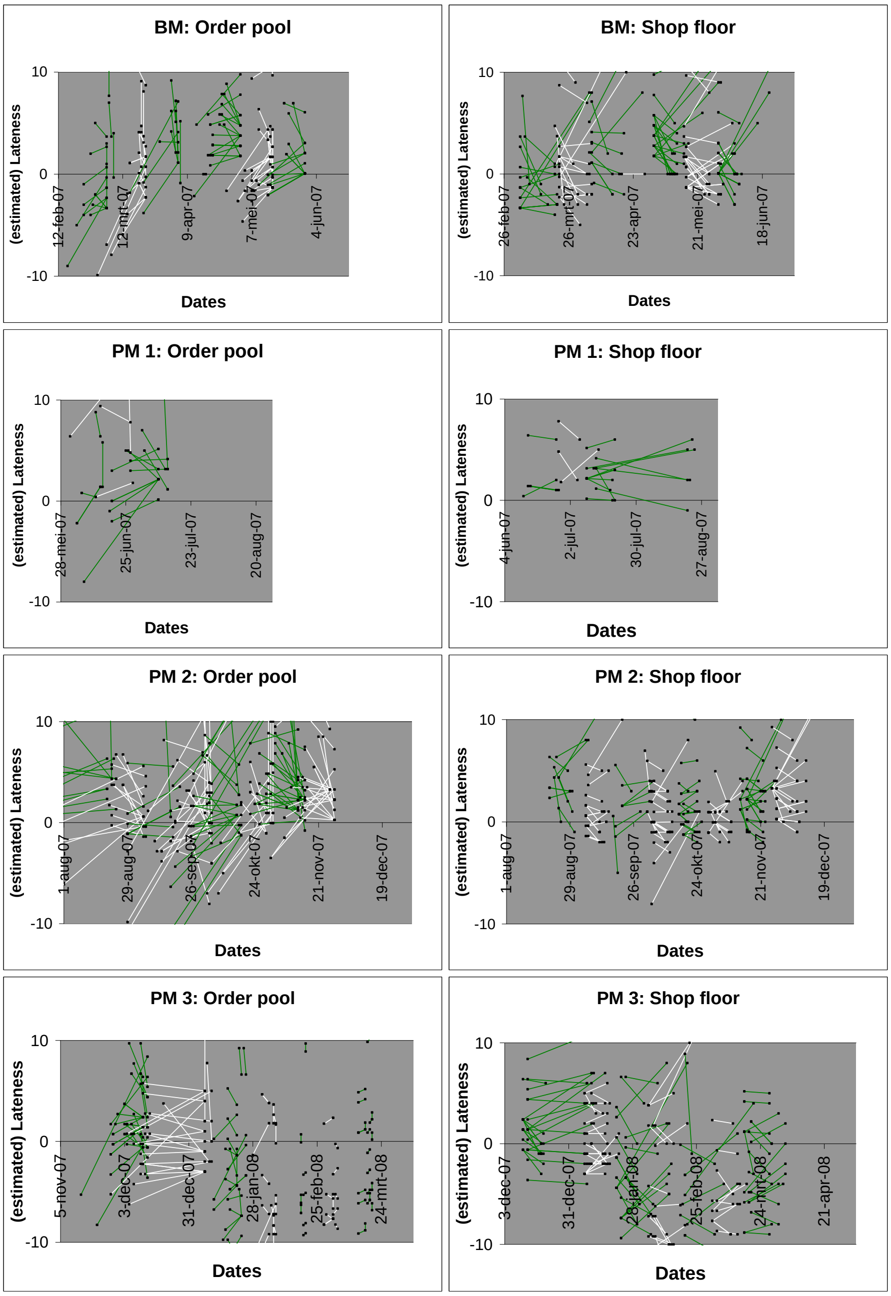

Figure 4.7. Order progress diagrams for each period.

4.5.1. Results - Analysis of the baseline measurement period

Analysis of average shop floor throughput times

The first part of the throughput diagram in Figure 4.6 shows the results for the baseline measurement period. The average level of WIP on the shop floor (5) is 4110 processing hours during this period. Figure 4.8 focuses in on this baseline measurement period.

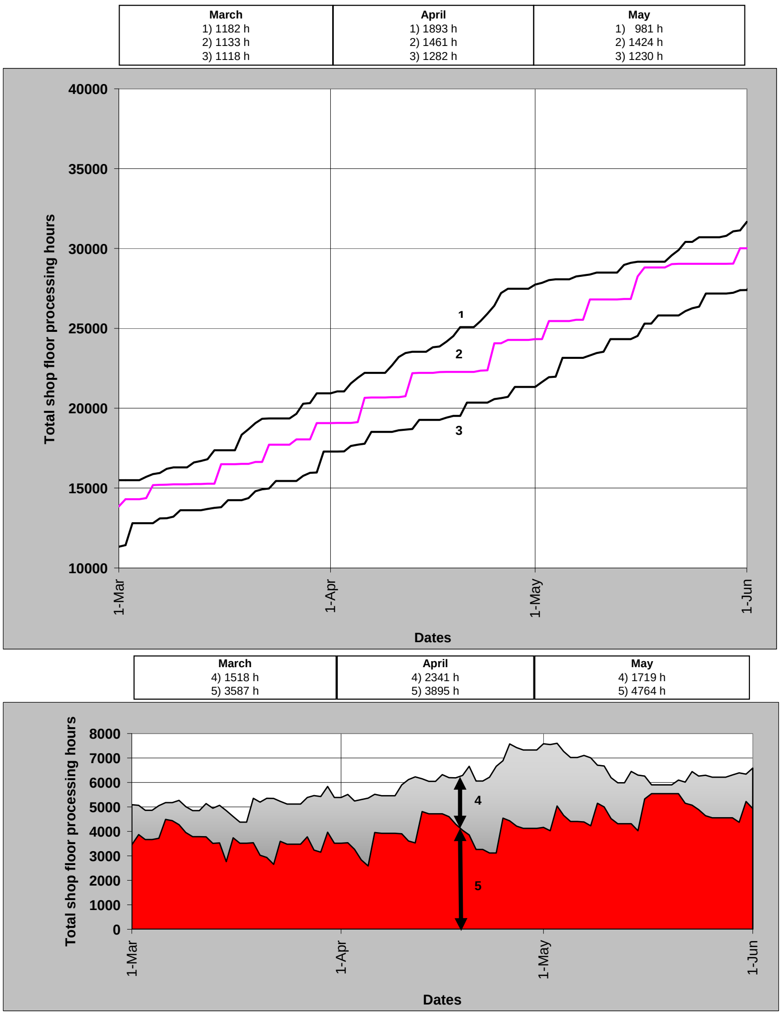

Figure 4.8. Throughput diagram of baseline measurement. 1) Order arrival; 2) Order release; 3) Order completion; 4) WIP in order pool; 5) WIP on shop floor. For curves 1 to 3, the average weekly input/output level (in hours) is specified per period. On top of the WIP graph, the average WIP levels (in hours) are specified per period.

Figure 4.8 shows that WIP on the shop floor varies considerably during the baseline measurement period, even when looking at monthly averages. For example, in March the WIP is on average 3587 processing hours, while in May the average WIP is 4764 processing hours. Since an average order takes 39 processing hours, 3587 processing hours reflects an average queue length of five orders per station (3587 processing hours/18 stations/39 processing hours per order). 4764 processing hours implies an average queue length of seven orders per station. Although normative conclusions are difficult to draw, the WIP level on the shop floor seems high and uncontrolled for a production situation with a dominant order flow pattern. These observed fluctuations are contrary to the aims of WLC.

The WIP on the shop floor (5) is dependent on both the order release rate (curve 2) and the order completion rate (curve 3). Figure 4.8 shows that the order release rate is irregular during the baseline period. The quantitative analysis reveals several causes underlying this irregularity. First, some of these fluctuations can be explained by the periodic, usually weekly, release pattern that causes the WIP on the shop floor to increase considerably each time there is a release. However, even after correcting for this effect the amount of hours released each week still fluctuates over time. Second, some of these fluctuations seem to be related to variations in the inflow of orders to the order pool. In particular, the strong surge of arriving orders in April, compared to March, results in a considerable increase in the number of orders released to the shop floor. As such, the order pool does not behave as an absorptive buffer, as intended by WLC.

To summarise, it can be concluded that the uncontrolled release of work to the shop floor hindered the control of average shop floor throughput times during the baseline measurement period.

Analysis of the variance in shop floor throughput times

The results, illustrated with the throughput diagrams, show that the average throughput time changed over time. Since the average is used as a reference point in the order progress diagram, a changing average results in upward and downward movements of the curves irrespective of the progress variations among the individual orders. To avoid these effects in the order progress diagrams, we compare the progress of each individual order with the average throughput times of all orders released in the same week. The resulting ‘normalised’ order progress diagrams for all periods are shown in Figure 4.7. Figure 4.9 provides an enlarged view of the baseline measurement period. The analysis, explained with the help of the diagrams, provides several insights.

Orders arrive in the order pool with large variations in estimated lateness, related to the tightness of the promised delivery date. Many of the orders with a significant estimated lateness on order arrival spend a shorter time in the order pool than orders that arrive with a low estimated lateness. This indicates that priority setting at release is often based on urgency. Further, a more detailed analysis of the data reveals that some urgent orders that arrive just after a regular release moment are not even retained until the next release moment, but are released almost immediately onto the shop floor. Despite the smoothing effect of these prioritising decisions, the resulting estimated lateness at release still differs among individual orders. This means that orders require different throughput times on the shop floor if they are to be completed on time.

On the shop floor, the progress of individual orders does indeed vary considerably. However, these variations in progress do not seem to all contribute to a reduction in the variation in the lateness of individual orders. This is partly a result of small disturbances such as inadequate priority decisions. Further, the analysis revealed both urgent and less urgent orders with extreme delays. In the workshop after the baseline measurement, a discussion with the production manager revealed that many of the orders in the latter group were ones with missing materials. As missing materials were generally only detected during production, the progress of the corresponding orders was disrupted on the shop floor. It is clear that such an approach does not provide a good basis for controlling the throughput times of orders on the shop floor.

Whereas the analysis of average throughput times showed that the applied release decisions resulted in fluctuating average throughput times on the shop floor, this subsection has shown that additional differences in the progress of individual orders are caused by urgency differences at release, material shortages and small progress disturbances on the shop floor (for example, caused by poor priority decisions). These baseline measurement results are used as inputs for the next research step.

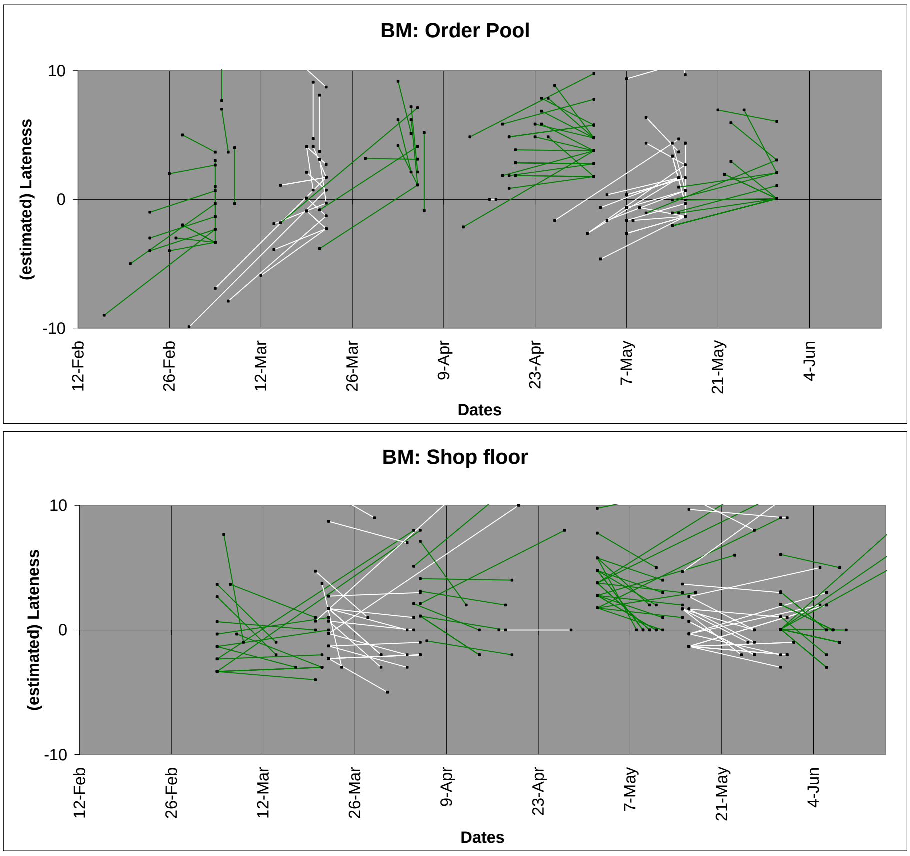

Figure 4.9. Order progress diagrams for baseline period.

4.5.2. Results - Incorporating WLC

Alongside presenting and discussing the results from the baseline measurement period, WLC was also introduced to the managerial staff attending the introductory workshop. The discussions in the workshop resulted in several ideas for improving the currently applied PPC decisions based on using the WLC mechanism. A couple of follow-up meetings were organised to elaborate on these ideas. These meetings resulted in agreements on the following major changes.

First, the main control decisions were to be shifted from taking place on the shop floor to taking place at release. The production planner agreed to aim for an improvement in the control of shop floor throughput times. As the average WIP level on the shop floor was viewed as rather high during the baseline period, this level could also be reduced to some extent. Next, central load balancing should be used to control the WIP remaining on the shop floor. However, central load balancing using explicit quantitative norms for each station was not realisable because information on estimated operation times was not available at the level of individual stations, and could not be obtained within a short period. Nevertheless, it was expected that a certain level of central load balancing could be achieved based on the production planner's knowledge of the production situation. For example, some smoothing could be accomplished by controlling the number of doors and windows released, as certain quantities were good indicators of potential high loadings at specific stations. An additional measure was that only those orders that could be fully completed would be released. This was important because even perfectly balanced queues would not lead to stable shop floor throughput times if orders had to wait on the shop floor for missing materials. To achieve this change, the production planner divided the order pool into two parts. The first part consisted of those orders for which all the raw materials were available and the second category contained orders with incomplete materials. The status of orders was regularly checked and orders were immediately transferred to the first category on the arrival of the missing material.

These measures were introduced at the beginning of June. The next section discusses the effects of these changes over time, with the post-change measurement period divided into three.

4.5.3. Results - Analysis of the post-change measurement period

The total post-change measurement period lasted approximately 10 months. We distinguish three phases within this period (see Figure 4.6). The first period was a transient one in which some of the WIP on the shop floor was being shifted to the order pool. In the second period, a more-or-less stationary situation was achieved in terms of the WIP on the shop floor. Finally, the third period reflects another transient phase, in which the effects of an unforeseen reduction in the number of arriving orders are observed. Analysing these successive periods in detail provides insights into the functioning of the implemented WLC measures in a practical dynamic setting.

Post-change Period 1 (June - July)

The first period of post-change measurements (PM 1) saw a considerable decrease in the WIP on the shop floor (indicated by line 5 in Figure 4.6). While the WIP on the shop floor was around 5000 hours at the end of May, this had reduced to around 3000 hours in mid-July. Note that the flat horizontal line starting mid-July represents the summer holiday when the company closed for four weeks.

Theoretically, a decrease in the WIP on the shop floor could be caused by either an increase in the order completion rate or a decrease in the release rate or some combination thereof. The average weekly rates for each month are provided at the top of Figure 4.10 and show that both changes did occur here. A more detailed analysis and discussion with the production manager revealed that extra overtime was being used in this period and a number of extra employees were hired to increased completions. In addition, a decrease in the weekly number of hours of work released also explains a considerable part of the reduction in WIP. It should be noted that the planner in this company was not used to aiming for a limited WIP on the shop floor. Intuitively, we would therefore expect him to find it challenging to reduce the release rate in a period where the order completion rate is increasing. Nevertheless, both the production planner and the production manager were enthusiastic over the new approach from the start. One of the reasons was the supportive reactions received from the sawing department, where workers observed that they now spent less time searching for non-available materials.

Logically, a reduction in the release rate leads to an increase of WIP in the order pool. Linked to this, it can be noted that, in terms of average weekly rates, the reduction in the release rate in June is coupled to a strong increase in the order arrival rate (curve 1). The left part of Figure 4.10 shows that the combined changes result in an increase in the share of the total WIP in the order pool from 27 % in May to 47% in July. Thus, a considerable buffer ahead of the shop floor was actually created over this period.

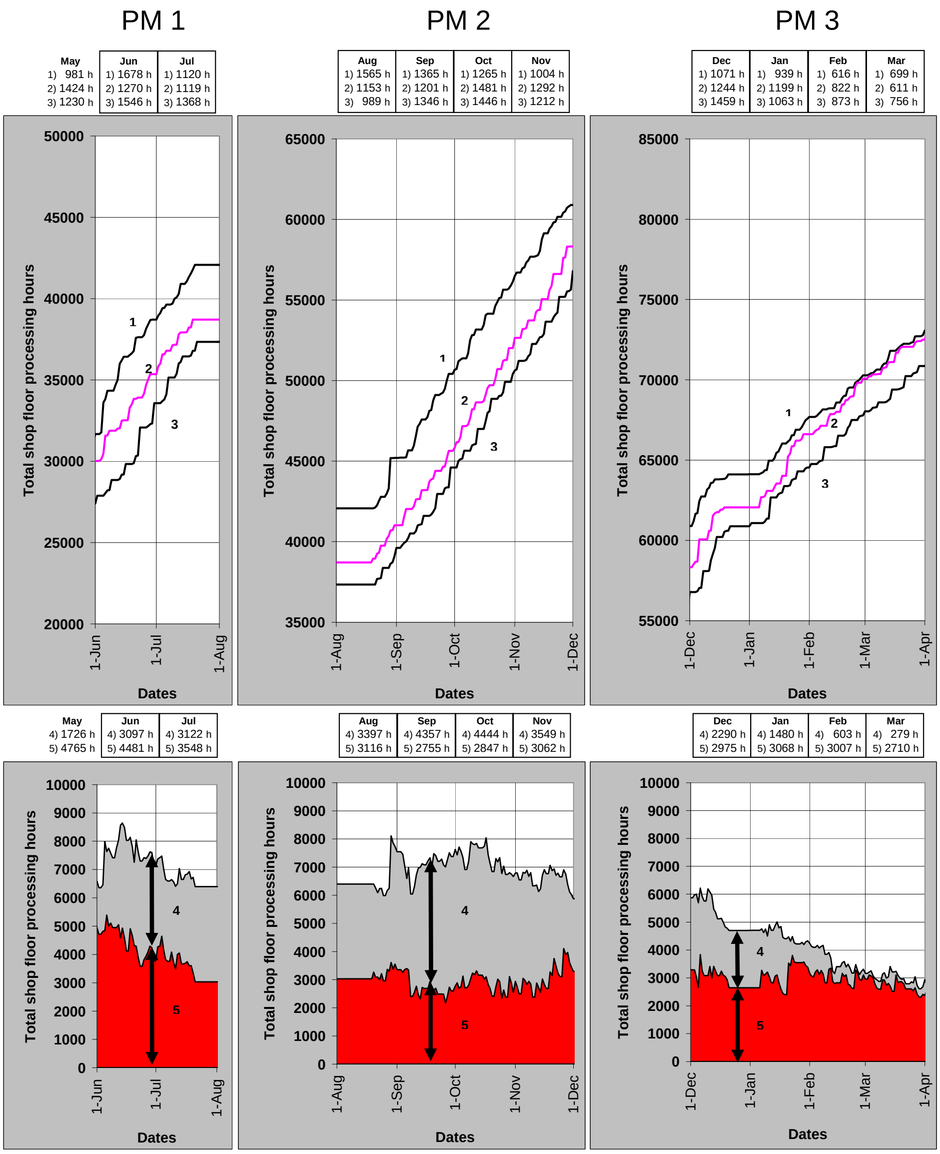

Figure 4.10.Throughput diagrams for each post-change period. 1) Order arrival; 2) Order release; 3) Order completion; 4) WIP in order pool; 5) WIP on shop floor. For curves 1 to 3, the average weekly input/output level (in hours) is specified per period. On top of the WIP graph, the average WIP levels (in hours) are specified per period.

Post-change Period 2 (August – November)

In the second period of the post-change measurements (PM 2), the reduction in the WIP on the shop floor ceased. The analysis regarding the WIP on the shop floor shows that, from that time on, the WIP has been relatively stable compared to the baseline period. The average weekly input and output rates at the top of Figure 4.6 show that the stable WIP is a result of the release rate being synchronised with the order completion rate. It should be noted, however, that even in this stable period the release rate is not always equal to the order completion rate in every month (see Figure 4.10), inevitably some hours of work released in one month will not be completed until the next month. The analysis regarding the order pool shows, in contrast to the baseline period, that fluctuations in the inflow number of hours to the order pool no longer results in a change to the number of hours released. As such, the order pool does act as a buffer for fluctuations in order arrivals, which enables average shop floor throughput times to remain stable over time.

The effects of the changes on the progress of individual orders are illustrated in the order progress diagrams for the second post-change measurement period (PM 2) in Figure 4.7. First, it can be observed that urgency-based priority decisions seem now to be taken within the order pool; often urgent orders are accelerated and non-urgent orders held back. Second, the effects of the increase in average throughput time in the order pool compared to the baseline period on individual order progress can be clearly seen. As the order pool contains more orders, priority setting can be more effective. The result is that the company is able to achieve a lower variance in the estimated lateness *at release* than in the previous period. This is a significant achievement because the variance of the estimated lateness *at the time of order arrival* had increased since the baseline period.

In addition to changes regarding order progress in the order pool, changes on the shop floor can also be detected. The results show that the variance of shop floor throughput times is somewhat reduced. The decreased level of WIP on the shop floor reduces the *opportunities* to change priorities between orders on the shop floor. Further, the shop floor variance has probably been affected by a reduced *need* for urgency-based priority setting given that the variance of the estimated lateness at release has been reduced. However, only tentative conclusions can be drawn on this point, as a reduction in variance could be anticipated from the reduction in average throughput times. Finally, it should be noted that, despite the measures taken, some extreme outliers can still be observed in this period and these can be attributed to material shortages on the shop floor. Nevertheless, the controlled release had enabled an improvement in the control of shop floor progress for most individual orders in this period.

Reporting these first positive results to the company's personnel during this period initiated some additional improvement actions to reduce the remaining variability in shop floor throughput times. Some workers at specific work stations were assigned the task of producing special parts in parallel to the main order flow, a so-called shunting approach. This change was only partly implemented at the end of this reporting period and so these changes are only marginally reflected in the data.

The observations of company personnel reflected the conclusions drawn from the measurements in this period. The production planner was rather enthusiastic, because he could easily see the benefits of having a more transparent shop floor, such as less effort being expended on searching for materials. He also became more confident in the predicted shop floor throughput times of orders, and started to notify the sales department when new orders were released to the shop floor. This information supported the sales department in updating customers about the planned delivery of their orders.

The results from the first two post-change measurement periods show that the main aim in applying WLC, namely improving the control of shop floor throughput times, was achieved. A logical next step would have been to improve the control of the WIP in the order pool. Controlling the WIP both in the order pool and on the shop floor would contribute to controlling overall throughput times from order arrival until order completion. However, circumstances changed in the last post-change measurement period as will be discussed in the next subsection below.

Post-change Period 3 (December – March)

The final measurement period can be characterised by a steady decrease in the total WIP (sum of 4 and 5 in Figure 4.10). Whereas, in the middle of November, total WIP was around 7000 hours, this had decreased to around 3000 hours by March. Theoretically, a decrease in the total WIP could be the result of a decrease in the order arrival rate and/or an increase in the order completion rate. In this instance, the main cause was a considerable reduction in the number of incoming orders. Despite this reduction, the WIP on the shop floor remained relatively stable during this third post-change measurement period (PM 3). A more detailed analysis of the WIP on the shop floor reveals that the average weekly order completion rate was reduced during this period (see top of Figure 4.10). This decrease was largely achieved by reducing the size of the workforce in the first weeks of January. The decrease in work output was accompanied by a reduction in the number of hours of work released to the shop floor. Comparable reductions in order completion and released hours results in the WIP on the shop floor remaining relatively stable. However, as could be expected from Little's result (Little, 1961), an unchanged WIP on the shop floor combined with a decrease in order completion rate results in an increase in average throughput times.

The effects of this on the progress of individual orders can clearly be seen in the order progress diagrams for this period in Figure 4.7. These diagrams show that incoming orders are released almost immediately to the shop floor. This indicates that release decisions are hardly used to prioritise urgent orders in this period. This is a result of the sharply reduced number of orders in the order pool; with so few orders there is hardly any scope to prioritise. As a consequence, orders are, by and large, released on a first-in-first-out principle, and this leads to a requirement for different throughput times on the shop floor to ensure they are completed on time.

Once on the shop floor, both the increased average throughput times and the lack of priority setting at release means that many orders have to be speeded up considerably to meet delivery dates, whereas others can be put to one side. Thus, it can be concluded that, despite the continuing limits on and control over WIP, the improved control of the progress of individual orders seen in the previous monitoring period (PM 2) could not be maintained in this final period.

4.6. Implications for WLC theory

The results of this study raise new perspectives on the functioning of WLC in a dynamic setting. First, the results reveal that a period of time is needed to reduce the WIP on the shop floor (PM 1). During this period, the company simultaneously increased capacity and decreased the number of hours of work released. The combined effect was a considerable decrease in WIP in a relatively short period. A few months later the company was confronted with a sharp decrease in demand, which could not be buffered using the order pool (PM 3). Again the company adjusted both capacity and the amount of work released. These adjustments led to both an increase in average throughput time and an increase in the variance of throughput times. The dynamics in the first period were the result of a change imposed and desired by the company itself, whereas the changes imposed in PM 3 were driven by externally created dynamics. However, both circumstances posed similar questions within the context of WLC, questions on the timing of adjustments regarding control decisions and the magnitude of these adjustments.

First, the importance of determining *when* to react in a dynamic setting should be stressed. Although timing is trivial when introducing WLC, reacting to observed external changes when using WLC is a more delicate affair. When an external change is first observed, it is often unknown whether it is a structural change that requires adjustments to the target WIP level. If a company fails to react to a structural change, performance may suffer. Conversely, overreaction may lead to system nervousness (Land and Gaalman, 1996). There is little literature on this timing issue as most WLC studies apply a WIP target level which is fixed over time. These target levels are often referred to as workload norms (Land and Gaalman, 1996) or workload bounds (Bergamaschi et al., 1997). One of the few methods which deals with adjustments to target levels over time is that proposed by Zäpfel and Missbauer (1993), in which target levels are continuously recalculated. However, this method might be too complex to implement in a practical setting and could easily lead to nervousness. Therefore, more research is needed to provide practical guidelines on the timing aspects of WLC in a dynamic situation.

Second, it is important to decide *which* control decisions should be adjusted in order to achieve a transition to a new situation. Many approaches can be applied in theory to achieve a transition. In the company studied, the reaction was to change both capacity and release simultaneously, while a transition could also have been achieved by changing only one of these factors or by changing them sequentially. More research is required to gain insights into the effects of these alternatives on performance.

Third, in addition to decisions on the timing and choice of which specific controls to change, it is important to decide on the *magnitude* of changes. This is a complicated issue as even target levels are often difficult to define in advance. As a consequence, a trial-and-error approach is often used in practice (Silva, 2009). However, this approach assumes a stationary situation, which contradicts the dynamics generally observed in practice. This situation is especially prevalent in the MTO order industry. The dynamics of the incoming order flow to which these companies are subjected have a direct impact on the workloads, in contrast to the make-to-stock situation where other instruments can be used to handle these dynamics. More research is required to provide insights into optimal WIP levels under a range of circumstances in order to provide practical guidelines for companies in MTO industries.

These three dynamic issues, derived from an empirical study, provide some important areas for further research on WLC. Greater insight into these issues is required if WLC is to be used effectively in a dynamic practical situation.

4.7. Conclusions

This research aims to add to the very limited empirical evidence on the effects of applying WLC in a practical dynamic setting. For this purpose, performance changes have been analysed over time and linked to the intended functioning of WLC. The empirical findings show that the company, while using a WLC approach, had to deal with an environment which was much more dynamic than is generally assumed in simulation studies. Two types of dynamics have been distinguished that WLC has to be able to cope with: dynamics imposed by the company itself and externally-driven dynamics. The findings show that, despite the use of WLC, if the dynamics are not well considered, throughput times can become uncontrolled over time. This can be explained from the fact that the perceived functioning of WLC is strongly based on insights from stationary simulation models. While an order pool in front of release does provide a sufficient buffer in a stationary setting, its limitations in a dynamic setting were exposed in this study. While the research on WLC in stationary settings has tended to focus on parameter choices, this study shows that it is questions on timing and magnitude of control adjustments that need to be answered if WLC is to be used effectively in a dynamic setting.

The results of the research, highlighted in this paper, put a new perspective on the WLC paradox that had been one of the triggers for this research. The results show that the performance achieved using WLC varies over time, and that this is, at least to some extent, determined by environmental changes confronting a company. These dynamics tend not to be fully taken into account when performance changes are evaluated in both simulation studies and empirical studies. This study raises the question as to what extent these dynamics can explain the differences in WLC performance found between simulations and empirical studies; an interesting starting point for future research.

Appendix

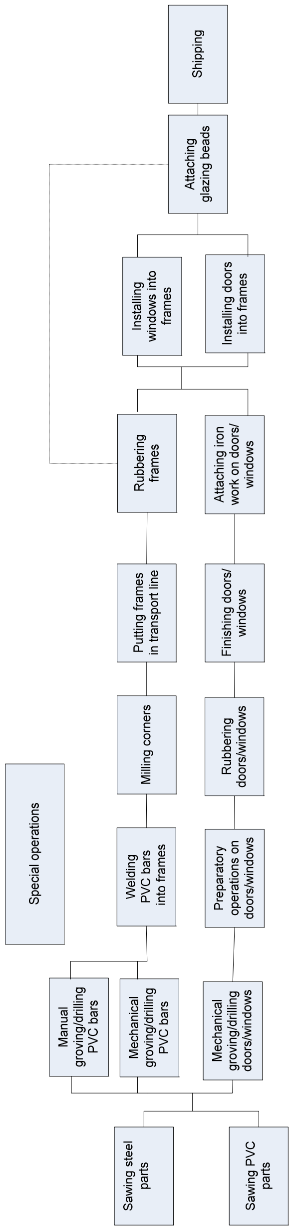

Figure 4.11. Operations on the shop floor.