University of Groningen

Load-Oriented Order Release (LOOR) revisited

Yan, Haoyun; Stevenson, Mark; Hendry, Linda C.; Land, Martin J.

Published in:
Production Planning & Control

DOI:
[10.1080/09537287.2016.1183831](https://doi.org/10.1080/09537287.2016.1183831)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2016

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Yan, H., Stevenson, M., Hendry, L. C., & Land, M. J. (2016). Load-Oriented Order Release (LOOR) revisited: bringing it back to the state of the art. *Production Planning & Control*, 27(13), 1078-1091. <https://doi.org/10.1080/09537287.2016.1183831>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Load-Oriented Order Release (LOOR) revisited: bringing it back to the state of the art

Haoyun Yan^a, Mark Stevenson^b, Linda C. Hendry^b and Martin J. Land^c

^aDepartment of Industrial Engineering, Shanghai Dian Ji University, Shanghai, P.R. China; ^bDepartment of Management Science, Lancaster University Management School, Lancaster, UK; ^cFaculty of Economics and Business, Department of Operations, University of Groningen, Groningen, The Netherlands

ABSTRACT

In the workload control literature, the Load-Oriented Order Release (LOOR) approach has been neglected since its robustness was questioned at the end of the 1990s. This paper revisits LOOR and evaluates whether its performance can be improved in two ways. First, an intermediate pull release mechanism is added to avoid starvation between periodic release events. This mechanism was recently shown to be effective at improving the performance of a state-of-the-art release method known as LUMS COR. Second, an integer linear programming model is used to manage the trade-off between the timing and load balancing functions of order release. The two refinements are assessed using simulations of different shop configurations, which allow us to evaluate robustness. Results demonstrate that the refinements contribute to improving the performance of LOOR such that it can even outperform LUMS COR. Perhaps counter-intuitively, putting more emphasis on load balancing than on the urgency of individual orders is shown to lead to a lower percentage of tardy orders. Overall, the improvements mean that concerns about LOOR's robustness are no longer valid – it now appears suitable for a wide range of shops found in practice.

ARTICLE HISTORY

Received 8 July 2015
Accepted 11 April 2016

KEYWORDS

Workload control (WLC);
intermediate pull release;
load balancing; job shop;
simulation

1. Introduction

Workload control (WLC) is a production planning and control concept of particular relevance to high-variety make-to-order shops (e.g. Stevenson, Hendry, and Kingsman 2005). A key feature of WLC is the use of a pre-shop pool and order release mechanism. Orders are held back from the shop floor in a pool and released in time to meet due dates while limiting work-in-process on the shop floor. There are three broad sets of release methods in the WLC literature (see Land and Gaalman 1996; Thürer, Stevenson, and Silva 2011) that originate from research conducted at the University of Hannover (e.g. Bechte 1988), Lancaster University (e.g. Tatsiopoulos and Kingsman 1983), and the Eindhoven University of Technology (e.g. Bertrand and Wortmann 1981). Most recent attention has been on variants of the methods from Lancaster and Eindhoven, including the studies by Thürer et al. (2012, 2014). In contrast, the method developed at Hannover, known as Load-Oriented Order Release (LOOR), has been largely neglected since its robustness was criticised. In particular, Land and Gaalman (1998) and Oosterman, Land, and Gaalman (2000) questioned the effectiveness of LOOR in shops with a dominant flow direction and when combined with a due date-oriented dispatching rule. LOOR does however have some advantages over other WLC release methods as it has a more advanced approach to load balancing that looks at how orders move through the shop over time. Yet, the procedure it applies for selecting orders for release is still considered to be

myopic – it does not look ahead to slightly less urgent orders that would better balance the load (Land and Gaalman 1996). In this respect, it does not optimise the trade-off between the timing and load balancing functions of order release. Optimisation is in general a neglected approach in the WLC literature since the seminal work of Irastorza and Deane (1974). In this study, we combine the two neglected areas (of LOOR and optimisation) together. We revisit and improve LOOR, including through the use of optimisation in the form of an integer linear programming model.

As recent research has ignored LOOR, this method – which only releases orders periodically – has not taken advantage of some of the latest developments in the WLC literature that have culminated in the LUMS COR (Lancaster University Management School Corrected Order Release) method presented by Thürer et al. (2012). The authors' results for LUMS COR show that combining periodic with continuous order release in the form of an intermediate pull release mechanism is particularly important as it avoids premature work centre idleness before the next release opportunity.

LOOR has been a key release method in the WLC literature, and it has been incorporated in many Enterprise Resource Planning (ERP) systems (Breithaupt, Land, and Nyhuis 2002). Thus it is important to consider whether its performance can be improved. In the light of the above, two possible search directions for improving the performance of LOOR are explored in this paper

using simulation. First, we evaluate the performance effect of adding an intermediate pull release mechanism. Second, we evaluate a new integer linear programming model that overcomes the myopic nature of the order selection procedure at release, which should allow for an improved trade-off between timing and load balancing. We therefore start with the following two research questions:

- (1) Can adding an intermediate pull release mechanism to LOOR improve its performance?
- (2) Can the performance of LOOR be enhanced by improving the trade-off between the timing and load balancing functions of order release?

Performance will be compared throughout the paper with both the original LOOR approach and with LUMS COR.

The remainder of this paper is organised as follows. Section 2 provides the theoretical background to the study, including an overview of LOOR and the refinements being made to the approach. Section 3 then presents the experimental design before the results are presented in Section 4. Section 5 provides an analysis of the results before a conclusion is provided in Section 6, which includes implications for practice and future research.

2. Theoretical background

Section 2.1 provides a brief history of the theory behind LOOR before Section 2.2 presents a simple improvement to LOOR that was suggested in the literature but has not previously been evaluated. Section 2.3 then reviews recent developments in the WLC literature; analyses the strengths of LUMS COR; and proposes the first refinement to LOOR (referred to as LOOR+). Finally, Section 2.4 analyses two basic functions of WLC release methods – timing and load balancing – before proposing the second refinement to LOOR based on integer linear programming (ILOOR). We also combine the two refinements to form ILOOR+.

2.1. Load-Oriented Order Release

Load-Oriented Manufacturing Control was developed as an integrated production control approach at the University of Hannover in the 1980s and 1990s (e.g. Bechte 1988, 1994; Wiendahl, Glässner, and Petermann 1992; Wiendahl 1995). The integrated approach has three control levels: Load-Oriented Order Entry, LOOR, and shop floor dispatching. LOOR, which is our focus, is at the heart of Load-Oriented Manufacturing Control and regulates the release of orders to the shop floor.

LOOR is arguably more straightforward than other load-based order release methods in that it focuses on controlling the direct load – measured by the sum of processing times – that is available for processing at each work centre. Orders are released at regular, periodic intervals, referred to as the release period (Oosterman, Land, and Gaalman 2000). At each release event (or periodic interval), a rough estimation is made of the input that can be expected at a work centre from work that is: (i) newly released; and, (ii) currently still upstream. Land and Gaalman (1996) showed how norm setting for LOOR is based on a balance equation – see Equation (1) below – that holds for each release period and each work centre w .

$$D_w^B + I_w = D_w^E + O_w \quad (1)$$

with D_w^B : the direct load of w at the *beginning* of the release period, after the release decision, I_w : the input to the direct load of w during the release period, D_w^E : the direct load of w at the *end* of the release period, before the next release decision, O_w : the output of w during the release period.

A target workload norm level λ_w is specified for the right-hand side of Equation (1), which consists of two components: (i) a desired remaining direct load buffer before the next release decision, specified by the direct load norm δ_w ; and, (ii) the planned output level, specified as ω_w . The periodic release decision should then bring the left-hand side of Equation (1), i.e. the direct load D_w^B plus the input I_w towards the norm level λ_w . A load conversion algorithm is used in LOOR to make a rough estimate of the inputs based on the assumption that a fraction ω_w/λ_w will, on average, proceed to the next work centre during the release period (Breithaupt, Land, and Nyhuis 2002). As such, for every order j that will be released or that is still upstream of work centre w , a contribution x_{wj} to the estimated inputs is specified by Equation (2a).

$$x_{wj} = p_{wj} \prod_{i \in U_{wj}} \frac{\omega_i}{\lambda_i} \quad (2a)$$

for all w to be visited after the next operation of order j

with p_{wj} : the processing time of order j at work centre w , U_{wj} : the set of upstream work centres that order j has to pass before arriving at w .

If w is the imminent work centre in the routing of the order, the contribution relates to the direct load D_w^B . In that case, the contribution x_{wj} to the load is equal to its full operation processing time p_{wj} as given by Equation (2b).

$$x_{wj} = p_{wj} \quad (2b)$$

for the imminent work centre w in the routing of order j

The release procedure considers the orders j for release in sequence according to the earliest planned release date τ_j^R . All orders with a planned release date that falls within a specified time limit ϑ are candidates for release. An order is only released if its contributions x_{wj} will not cause D_w^B plus the estimated value of the future inputs I_w to exceed the workload norm λ_w of any work centre w in its routing. Otherwise, it has to wait in the order pool until at least the next release decision and the order with the next planned release date will be considered. Although the LOOR approach originally allowed one order to be released that would take a norm above its maximum level – otherwise a norm may never be completely filled – this has since been shown to be detrimental to performance (Oosterman, Land, and Gaalman 2000). The planned release date τ_j^R of an order j is determined according to Equation (3) by backward scheduling from its due date τ_j^* , allowing for a planned lead time γ_w for each work centre w in the routing R_j of the order.

$$\tau_j^R = \tau_j^* - \sum_{w \in R_j} \gamma_w \quad (3)$$

Using the time limit ϑ would restrict the set of orders that can be selected for release, reducing the potential for load balancing. Therefore, Land (2006) proposed the use of an infinite time

limit unless there are specific circumstances that prohibit this, e.g. high early completion costs.

After release, order progress on the shop floor is controlled by a priority dispatching rule. Bechte (1988) suggested combining LOOR with a simple first-come-first-served (FCFS) rule. But it was acknowledged that under certain conditions, like machine breakdowns or rush orders, due date-oriented dispatching rules can also be applied. Land (2004) however found that even due date-oriented dispatching rules will lead to the best performance of LOOR being at infinite workload norms – which is undesirable from the perspective of workload control – while recent empirical evidence calls for the use of more powerful priority dispatching rules than FCFS in practice (e.g. Soepenbergh, Land, and Gaalman 2012). Thürier et al. (2012) showed that other improved release methods can achieve their best performance with restricted workloads, including when combined with the Planned operation Start Time (PST) dispatching rule. In this paper, we will therefore examine the performance of LOOR and its variants with the PST rule to see if it too can achieve its best performance at lower (restricted) workload levels.

2.2. Ensuring robustness using a look-ahead limit in load calculations

Simulation results have shown that LOOR performs particularly well in a job shop environment (e.g. Land and Gaalman 1998; Cigolini and Portioli-Staudacher 2002). But when the flow is directed, such as in the general or pure flow shop, Oosterman, Land, and Gaalman (2000) found that LOOR is outperformed by aggregate load methods and concluded that this was because focusing on the direct load can create undesirable (cyclic) effects when a shop has typical upstream and downstream work centres. Breithaupt, Land, and Nyhuis (2002) later suggested these effects can be avoided in practice by only including the next four routing steps in LOOR's load calculations, i.e. Equation (2a). However, the influence on robustness of limiting how far LOOR looks ahead (a look-ahead limit) has not been evaluated. This study will be the first to incorporate and evaluate this simple solution in a simulation study (for LOOR and the refined versions of LOOR we present below).

2.3. Improving LOOR via intermediate pull release (LOOR+)

One of the dimensions used by Bergamaschi et al. (1997) to classify WLC order release methods was according to *when* the release decision takes place: periodically or continuously. Periodic release methods, including LOOR, release orders at fixed, regular intervals. Meanwhile, continuous release methods lead to releases occurring at any time, usually triggered by an event like the workload falling to a certain level. Periodic release methods have received much more research attention than continuous release methods and are more straightforward to implement in practice. For example, they do not rely on a continuous flow of information on order progress back from the shop floor (Bergamaschi et al. 1997; Thürier et al. 2012). But periodic release methods can result in premature idleness or starvation (Land and Gaalman 1998); and simulations have found that some

continuous release methods can outperform pure periodic release methods (Thürier et al. 2012).

LUMS COR uniquely combines periodic with continuous release. It uses the corrected aggregate load approach (e.g. Oosterman, Land, and Gaalman 2000) for periodic releases and a lower workload bound to pull orders onto the shop in-between periodic release decisions if a work centre is starving. More specifically, if there are no orders currently queuing or being processed at a work centre then all remaining orders in the pre-shop pool with this work centre as the first in their routing are considered for release. From this set, the order with the earliest planned release date is selected for release regardless of its load contribution. Although the periodic selection procedure is similar to the procedure of LOOR, load calculations in LUMS COR differ from those in LOOR. If w is the v -th work centre in the routing of order j , then the contribution x_{wj} of j to the load of w under LUMS COR is given by Equation (4).

$$x_{wj} = \frac{p_{wj}}{v} \quad (4)$$

This contribution will remain the same from the release of the order until the operation at work centre w has been completed, while the load contributions in LOOR change as an order moves through the shop towards work centre w .

The performance results for LUMS COR presented in Thürier et al. (2012) showed that combining periodic release with an intermediate (continuous) pull release mechanism can lead to substantial performance improvements as it can avoid premature idleness in the context of a (corrected) aggregate load order release method. The intermediate pull release mechanism has been shown to be effective not only in balanced but also in unbalanced shops (Fernandes, Land, and Carmo-Silva 2014). This mechanism may also enhance the performance of LOOR. Thus, the first way that we will attempt to improve LOOR is by including the same intermediate pull release mechanism – this approach is referred to as LOOR+.

2.4. Improving LOOR via load balancing using integer linear programming model (ILOOR)

2.4.1. Background: the trade-off between timing and load balancing

Land (2004) argued that WLC release methods should serve two functions: (i) to balance the workload by maintaining a constant direct load level at each work centre; and, (ii) to allow for timely release such that each order can meet its due date and expected flow time. On the one hand, load balancing may support the timing function – as it can reduce lead time variation, which means that the planned release dates of orders can be determined more accurately. On the other hand, the two functions can conflict, e.g. if non-urgent jobs that balance the load are selected ahead of more urgent jobs that do not balance the load (and have to wait in the pool). Below, we will seek to improve the order release selection procedure of LOOR while bearing in mind these two, potentially conflicting functions. Analytical research by Wein (1990) and Wein and Chevalier (1992) underlined the importance of giving careful consideration to these two functions of order release and advised, for example, that

order release should give priority to improving load balancing when there are only small differences between the urgency of jobs in the pool. However, the current LOOR procedure considers orders according to their urgency based on planned release dates; it does not consider how the degree of urgency differs between jobs.

2.4.2. Improving LOOR's order release selection procedure

LOOR's timing function depends principally on selecting orders for release in planned release date sequence, while its load balancing function depends on filling up the workload norms in this fixed sequence. Once an order fits within the norms, its release is not re-evaluated. This can be considered a greedy algorithm. As a result, the workload of some work centres could be far below their norm level because another work centre's workload has reached its limit, even though the planned release dates of selected orders may hardly differ from those of some non-selected orders. As an example, Figure 1(a) shows the workload of three work centres (A, B and C); and Figure 1(b) the workload contributions of three orders (1, 2 and 3) in the pool. Assume that Order 1's planned release date is slightly earlier than that of orders 2 and 3. According to LOOR's release procedure, Order 1 should be considered for release first; and since no norms are exceeded, it is released and its workload contributes to the relevant work centres. Work Centre A then refuses to accept any further orders. Therefore, orders 2 and 3 cannot be released as they both contribute to Work Centre A. Consequently, work centres B and C are left far below their norm level (Figure 1(c)) and two orders are delayed. Figure 1(d) shows an alternative where orders 2 and 3 are released and Order 1 is retained in the pool. A better overall workload balance is achieved and only one order is delayed, but this solution would not be found using LOOR's current procedure. This is exactly why the procedure was considered myopic by Land and Gaalman (1996).

Figure 1. Example workload implications of selecting different orders for release. (a) Workload of work centres. (b) Workload contribution of orders in the pool. (c) Workload of work centres if order 1 is released. (d) Workload of work centres if orders 2 and 3 are released.

From the above, it follows that there are opportunities to improve the order release procedure of LOOR and improve the trade-off decision between load balancing and timing. It is argued that the order selection procedure should consider the timing and load balancing functions simultaneously. One feasible way of achieving this is using an integer linear programming model to select orders. Irastorza and Deane (1974) introduced the use of linear programming to the order release literature, but the approach has received little attention since their contribution. Lödging (2013) recently suggested this may have been because linear programming models were difficult to solve with the computational power of the 1970s – clearly this is no longer the case.

Irastorza and Deane (1974) approached order release as an order selection problem where the objective function contained two parts: one to control deviations between the total planned and actual work-in-process on the shop floor, which represents the load balancing function; and the other to control the early release of orders, which represents the timing function. We will adopt a similar approach to Irastorza and Deane (1974) – but designed for the specific control target of LOOR – as described below (see Section 2.4.3).

2.4.3. The integer linear programming model (used in ILOOR and ILOOR+)

We will use a simple integer linear programming model to solve LOOR's order selection problem; the model is given by Equations (5) and (6) below.

$$\max z = \sum_{j=1}^n \left(\frac{1}{\max(s_j, 0.1 \cdot 0.95^{-s_j})} y_j \right) + \alpha \sum_{w=1}^W \left(\frac{L_w}{\lambda_w} \right) \quad (5)$$

$$\text{s.t. } L_w \leq \lambda_w \quad \forall w = 1 \dots W \quad (6)$$

The notation used in the model is as follows:

j : index of the orders considered for release $j := 1 \dots J$.

i : index of the orders on the shop floor $i := 1 \dots I$.

w : index of the work centres $w := 1 \dots W$.

s_j : slack of order j at time t of the release decision, i.e. $s_j = \tau_j^R - t$.

y_j : decision variable indicating whether order j is selected for release ($y_j = 1$) or not ($y_j = 0$).

L_w : load resulting after the release decision, i.e.

$L_w = \sum_i x_{wi} + \sum_j y_j x_{wj}$ with x -values calculated by load conversion, as specified in Equations (2a) and (2b).

λ_w : the workload norm level (see Section 2.1).

α : relative weight given to the load balancing function.

The first part of the objective function in Equation (5) represents the timing function, which gives priority to orders based on slack s_j with respect to the planned release date τ_j^R . To give priority to orders with small slack values, orders are selected to maximise the sum of $1/s_j$ values. The max-function is inserted to handle slacks approaching zero and negative slacks. It creates priorities that continue to increase moderately when the slack finally goes down to zero and becomes negative. This function has been adapted from Shafaei and Brunn (2000) and its parameters have been pre-tested to ensure that it results in the desired slack-based behaviour. The second part of the objective function helps to balance the load. The closer the calculated workload L_w is to

its norm level λ_w the higher its contribution to the maximisation objective. This balancing part of the model is weighted by a factor α compared to the timing part. In Section 5.1, we will explore how the value of α impacts performance, thereby providing an insight into how α should be set. The constraints in the model are given by Equation (6) and ensure that the workloads of work centres do not exceed their norm levels.

Solving the above model will generate the set of orders for release. This model can replace LOOR's original order selection procedure – and it becomes the second way in which we will attempt to improve LOOR, referred to as ILOOR (integer linear programming model-based LOOR). We will also combine the two improvement approaches, i.e. intermediate (continuous) pull release and integer linear programming, to create ILOOR+. The simulation model used to evaluate the refinements is outlined next in Section 3 before the results are presented in Section 4.

3. Simulation model

3.1. Shop and job characteristics

A simulation model has been developed using Arena 8.0. The model represents a shop with six work centres, where each is a single and unique capacity resource. Capacity is equal for all work centres and remains constant. Two shop configurations have been modelled: the pure job shop of Melnyk and Ragatz (1989) and the general flow shop of Oosterman, Land, and Gaalman (2000). Under both configurations, an order does not visit the same work centre twice and all work centres have an equal probability of being visited. Hence, the maximum number of operations per order is 6. Each operation requires one specific work centre and both routing and operation processing time characteristics are known upon order entry. The routing length is determined by drawing from a discrete uniform distribution [1, 6]. The routing sequence is completely random in the pure job shop. In the general flow shop, the set of work centres is random but the sequence is in order of increasing work centre index (i.e. from 1 through to 6).

To aid comparison, order characteristics are as in Thürier et al. (2012). Thus, processing times p_{wj} follow a truncated 2-Erlang distribution with a non-truncated mean of 1 time unit and a maximum of 4 time units for every work centre w . Orders arrive according to a Poisson process, with the arrival rate such that the utilisation level is 90%. To set order due dates τ_j^* , a random allowance uniformly distributed between 0 and 30, and a minimum value covering a minimum shop floor throughput time, are added to the order entry time. The minimum value corresponds to the maximum processing time (4 time units) for the maximum number of possible operations (6) plus another 4 time units to allow for some waiting time.

3.2. Order release methods

Six release methods are simulated: LOOR, the three refined variants of LOOR (i.e. LOOR+, ILOOR, ILOOR+), and two reference methods (LUMS COR and immediate release, IMM where no order release control is applied). The procedures and calculations for LOOR, its variants and LUMS COR are implemented as specified in Section 2. Under IMM, orders are released as soon as they arrive at the shop.

To aid comparison, the same parameters are used for LOOR and each of its variants wherever possible. As in Land (2006), the release period is set to 5 days and, based on the arguments presented in Section 2.1, an infinite time limit ϑ is used, which means that all orders in the pool are considered for release. A planned work centre lead time γ of 4 days is used in all experiments, which is consistent with the average realised work centre lead time across the most relevant range of workload norm levels considered (see Section 3.4 below). As suggested in Land (2006), we use a constant planned lead time γ when calculating planned release dates τ_j^R (see Equation (3)) as performance is not very sensitive to this parameter. As in previous LOOR studies (e.g. Bechte 1988; Oosterman, Land, and Gaalman 2000), orders in the pool are sequenced according to their planned release dates.

Instead of applying a Load-Oriented order selection procedure, ILOOR and ILOOR+ release orders according to the solution proposed by the integer linear programming model. The model is solved at each periodic release event using CPLEX software embedded in the simulation model. The relative weight of α in the objective function of the model is 1000 in the main experiments described in Section 4. Later, in Section 5.1, we will evaluate the sensitivity of performance to the value of α .

3.3. Shop floor dispatching rules

Dispatching follows the PST rule, as used in recent WLC studies, including Thürier et al. (2012). The planned start time τ_{wj}^S of the operation performed by work centre w is given by Equation (7). R_{wj} is the set of remaining operations when order j has arrived at work centre w , and k is an allowance for the waiting time per operation; all other variables were defined in Section 2.1.

$$\tau_{wj}^S = \tau_j^* - \sum_{i \in R_{wj}} (p_{ij} + k) \quad (7)$$

These PST calculations reflect backward scheduling from the order due date and allow for both the processing time and a waiting time k at each remaining operation. The waiting time allowance k is set to 2 time units, as in Thürier et al. (2012). All experiments were also conducted with FCFS dispatching, as used in many early studies on WLC. As the choice of dispatching rule between PST and FCFS did not affect the relative differences between the release methods, and the use of FCFS generally weakened performance overall, only the results with PST dispatching are presented.

3.4. Experimental design and performance measures

We compare the order release methods at eight different levels of norm tightness. As LOOR and its variants use different workload aggregations to LUMS COR, we use the average shop floor throughput time as an intermediate variable to compare the release methods at different tightness levels, as in Oosterman, Land, and Gaalman (2000), Land (2004), and Thürier et al. (2012). For LOOR and its variants, we use the following 8 direct load norm levels, $\delta_w := 2.5, 3, 4, 5, 6, 8, 10$ and 20 time units. The planned output component ω_w of the workload norm is 4.5 in all experiments based on a 90% utilisation level and a release period of 5 time units. For LUMS COR, we set the following 8 workload norm levels: 4, 5, 6, 7, 8.5, 10.5, 12.5 and 24.5 time

units. As shown by Equations (2) and (4) above, load calculations differ between LOOR and LUMS COR. But these norm levels generate roughly similar flow times on the shop floor, which aids comparison. To facilitate full comparability, norm tightness will be indicated by the mean realised flow time of orders, i.e. the mean time that orders spend on the shop floor after release. Hence, we will refer to equal norm tightness when two methods result in the same mean flow time. The key performance measures, as listed below, will be shown graphically (set against the mean flow time):

- the mean lead time, which sums the time in the order pool waiting for release and the flow time;
- the standard deviation of lateness; and,
- the percentage of orders delivered tardy.

An improved focus on load balancing will be reflected in a decreasing mean lead time, as better balancing speeds up orders on average. An improved focus on the timing function of order release will normally be reflected in a lower standard deviation of lateness. Finally, improving the trade-off between balancing and timing – combining a short mean lead time and a small standard deviation of lateness – should reduce the percentage of orders that are delivered tardy.

Each experiment consists of 50 runs and results are collected over 10,000 time units during each run. The warm-up period is set to 3000 time units to avoid start-up effects. These parameters are in line with previous studies that applied similar job shop and general flow shop models (e.g. Oosterman, Land, and Gaalman 2000; Land, 2004, 2006; Thürier et al. 2012). Further, to reduce variance, we have used the common random numbers technique. To comply with variance reduction, significance can be proven based on a paired *t*-test whenever we present differences between two experiments.

4. Results

The main simulation results are presented in the following three subsections. First, the effect of adding an intermediate pull release mechanism is evaluated in Section 4.1. The use of an integer linear programming approach to adjust the weight given to load balancing and timing within the order release selection procedure is then assessed in Section 4.2. Finally, we evaluate the performance of the two refinements combined in Section 4.3.

4.1. Adding an intermediate pull release mechanism: LOOR+

The results for LOOR, LOOR+, LUMS COR and IMM in the pure job shop and general flow shop are illustrated in Figure 2. Figure 2(a) shows the lead time performance of each release method plotted against the mean flow time. A curve is constructed for each release method except IMM, which is represented by a single point. Each point on a curve is the result of simulating a release method at a specific norm level; hence, each curve contains eight points. From Figure 2(a), we can see that the curves for LOOR+ and LUMS COR almost converge at the furthest right-hand point as the two approaches will perform similarly if release is not restricted by a workload norm. At the highest workload norm level, LOOR is located to the upper right of LOOR+. This is because

the intermediate pull release element in LOOR+ means orders do not have to wait until the next release period for release to an idle work centre. In contrast, the purely periodic nature of LOOR restricts its performance. As the norms get tighter, i.e. moving from right to left along the curves, the mean lead time of LOOR increases, which is a major drawback of LOOR that has previously been highlighted in the literature. For LOOR+ and LUMS COR, the mean lead time first decreases and then increases, which shows this drawback can in fact be overcome. These two curves are close to one another, and they are located far below the curve for LOOR. Both are superior to LOOR in terms of mean lead time performance, but LUMS COR performs slightly better than LOOR+. The single result for IMM is located to the lower right of LOOR and to the upper right of both LOOR+ and LUMS COR. This demonstrates that LOOR reduces the flow time on the shop floor compared to IMM but at the cost of an increase in the mean lead time. In contrast, LOOR+ and LUMS COR reduce the flow time and lead time simultaneously when compared to IMM.

While a shorter mean lead time indicates that LOOR+ and LUMS COR improve load balancing, or more specifically avoid premature idleness (Thürier et al. 2012), the standard deviation of lateness results in Figure 2(b) will indicate whether this has been achieved at the expense of the performance of the timing function. Figure 2(b) shows that, in general, timing worsens as the norms are tightened, i.e. when we move to markers at the left-hand side of the curves. However, LOOR+ and LUMS COR show improved timing when compared to LOOR at tight norm levels.

Figure 2(c) presents the final consequences for delivery performance in terms of the percentage tardy. The relative positioning of the curves is the same as in Figure 2(a); hence, LUMS COR performs slightly better than LOOR+ in terms of the percentage tardy, but both perform better than LOOR. The LOOR curve continues to slope downwards while the curves for LOOR+ and LUMS COR achieve a minimum percentage tardy before the highest workload norm level is reached. All three methods outperform IMM on delivery performance.

Finally, performance in the general flow shop is depicted in Figure 2(d)–(f). The LOOR curves have the same shapes as in the pure job shop (Figures 2(a)–(c)). In earlier studies (e.g. Oosterman, Land, and Gaalman 2000), the performance of LOOR deteriorated much more in the general flow shop – and this difference in performance confirms that using the adaptation proposed by Breithaupt, Land, and Nyhuis (2002), whereby the load calculation only considers the next 4 operations, is effective. Compared to the pure job shop results, the standard deviation of lateness (Figure 2(e)) across the three methods is more equal. The relative positioning of the LOOR+ curves for the lead time (Figure 2(d)) and the percentage tardy (Figure 2(f)) are comparable to those in the pure job shop, suggesting that these results are robust to a change in shop configuration.

Overall, we can conclude that adding an intermediate pull release element leads to an improvement in the performance of LOOR.

4.2. Integer linear programming to improve the timing/load balancing trade-off: ILOOR

The results for LOOR, ILOOR, LUMS COR and IMM in the pure job shop and general flow shop are illustrated in Figure 3. Figure 3(a)

Figure 2. Simulation results for LOOR+ vs. LOOR, LUMS COR and IMM. (a) Lead time performance in the pure job shop. (b) σ of Lateness performance in the pure job shop. (c) Percentage tardy performance in the pure job shop. (d) Lead time performance in the general flow shop. (e) σ of Lateness performance in the general flow shop. (f) Percentage Tardy performance in the general flow shop.

presents the lead time performance in the pure job shop. Here, the ILOOR curve lies below the LOOR curve but above the LUMS COR curve, which shows ILOOR is an improvement over LOOR but is inferior to LUMS COR on mean lead time performance. As we would expect, the performance of ILOOR converges on that

achieved by LOOR at the highest workload norm level, i.e. to the far right of the figure. Here, the influence of the workload norm is almost zero and the order release selection procedure has no influence as all orders are likely to be released to the shop floor. As the norm level is gradually tightened, ILOOR's mean lead time

Figure 3. Simulation results for ILOOR vs. LOOR, LUMSCOR and IMM. (a) Lead time performance in the pure job shop. (b) σ of Lateness performance in the pure job shop. (c) Percentage tardy performance in the pure job shop. (d) Lead time performance in the general flow shop. (e) σ of Lateness performance in the general flow shop. (f) Percentage tardy performance in the general flow shop.

first decreases and then begins to increase, which shows ILOOR also overcomes the drawback with LOOR in this respect.

The performance of the timing function of ILOOR is rather similar to that of LOOR at the current parameter setting. This is reflected in the comparable levels of the standard deviation of lateness in Figure 3(b). LUMSCOR however performs better in

terms of the standard deviation of lateness at the tight norm levels to the left of the curve. The influence of the weighting given to parameter α on the load balancing and timing objectives in ILOOR will be explored further in Section 5.

In Figure 3(c), we see that the shape of the ILOOR curve for the percentage tardy is rather similar to that of the mean lead

time in Figure 3(a). In addition, the curve is very close to that of LUMS COR, which suggests ILOOR performs well in terms of the percentage tardy but does not quite achieve the performance level of LUMS COR at tight norm levels.

The improvement achieved by ILOOR over LOOR is also robust to a change in shop configuration, as demonstrated by the general flow shop results presented in Figure 3(d)–(f). Overall, we can conclude that the improved trade-off between load balancing and timing through the use of integer linear programming leads to an improvement in the performance of LOOR. However, this trade-off depends on the role of α which deserves further analysis in Section 5.

4.3. Combining the two refinements: ILOOR+

Figure 4 compares the performance of ILOOR+, which combines the above two refinements, with LOOR, LOOR+, ILOOR, LUMS COR and IMM. Figure 4(a) depicts the mean lead time performance in the pure job shop. As expected, we see that LOOR+, ILOOR+ and LUMS COR converge at the highest workload norm level. We also see that the three curves are very close to one another across a large range of flow times. Hence, the three methods that incorporate intermediate pull release result in similar performance. However, by combining the two refinements, ILOOR+ seems to perform slightly better than LOOR+ and LUMS COR in terms of the mean lead time (Figure 4(a)). However, the dispersion of lateness among jobs is slightly higher for ILOOR+, as indicated by the higher standard deviation of lateness for ILOOR+ in Figure 4(b). As a consequence, the final influence of ILOOR+ on the percentage tardy (Figure 4(c)) is rather similar to that of LOOR+ and ILOOR+. The curves in Figure 4(c) cross each other, which implies that the best-performing approach depends on the workload norm level applied, as reflected in the mean flow time on the shop floor. This performance is also sustained in the general flow shop (Figures 4(d)–(f)).

Overall, we can conclude that combining the two refinements leads to a marginal performance improvement, although ILOOR+ can slightly outperform LUMS COR at certain workload levels. The next section will seek to obtain a better understanding of the underlying mechanisms that lead to the performance results described. Section 5.1 analyses how the relative weights given to load balancing and timing impact the performance of ILOOR and ILOOR+ before Section 5.2 focuses on how adding an intermediate pull release mechanism improves the performance of periodic release methods.

5. Analysis of results

5.1. Sensitivity to the weight given to load balancing and timing in ILOOR and ILOOR+

The performance of ILOOR and ILOOR+ will be affected by the relative weight given to load balancing and timing within the objective function of the integer linear programming model, which is represented by the value of α . We have therefore conducted a sensitivity analysis to establish how the value of α , which was set to 1000 (10^3) in the main experiments in Section 4, affects performance. First, we set the value of α at two extremes, i.e. $\alpha = 10^{-6}$ and $\alpha = 10^6$. The former gives full priority to the

timing function while the latter gives full priority to load balancing. Finite and non-zero values of α are preferred to handle ties in the decision concerning which order to select when only one function is considered.

The simulation results in the pure job shop and the general flow shop for ILOOR and ILOOR+ are shown in Figure 5 together with the results for $\alpha = 10^3$ (as used in Section 4). For reference, we also present the curves for LOOR and LUMS COR.

From Figure 5(a), we can observe that the mean lead time of ILOOR (solid triangular markers) in the pure job shop is sensitive to α while the performance of ILOOR+ (open triangular markers) is less sensitive. This conclusion also holds for the general flow shop (Figure 5(d)). We can see in Figure 5(a) and (d) that the curves for ILOOR and ILOOR+ with $\alpha = 10^6$ (dotted curves) are just slightly below those with $\alpha = 10^3$ (solid curves), which shows that giving further emphasis to load balancing leads to only a marginal improvement in lead time performance. When priority is given to the timing function ($\alpha = 10^{-6}$, dashed curves), we see that the mean lead time of ILOOR and ILOOR+ deteriorates, but both still outperform LOOR (solid curve with square markers). Figure 5(b) and (e) shows that a strong emphasis on load balancing by using $\alpha = 10^6$ (dotted curves) increases the standard deviation of lateness slightly, while a strong emphasis on timing by using $\alpha = 10^{-6}$ (dashed curves) creates a weak advantage in the standard deviation of lateness. As the latter influences are relatively small, the percentage tardy (Figure 5(c) and (f)) is mainly affected by the changes in lead time performance (Figures 5(a) and (d)) that are due to improved load balancing.

The results for the percentage tardy may suggest that a complete focus on load balancing by setting $\alpha = 10^6$ is more appropriate than setting $\alpha = 10^3$, but this might cause a small number of orders to suffer from extreme lateness given the increasing standard deviation of lateness. To help in selecting a suitable value for α , we conducted additional simulation experiments where we fixed the direct load norm δ_w of ILOOR and varied the value of α stepwise from 10^{-6} to 10^6 . We set the direct load norm $\delta_w = 4$ because ILOOR performs well at this norm level (see Figure 3(a)–(c), third marker on the curves). The results of these additional simulations are presented in Figure 6 for the pure job shop and general flow shop, with the pattern of results being roughly similar across the two shop configurations. Figure 6(a) and (b) uses a logarithmic scale by setting the value of $\log_{10}(\alpha)$ on the horizontal axis. Moving from left to right in each of the graphs demonstrates how shifting the focus of the release selection procedure from timing to load balancing improves both lead time and percentage tardy performance. Below a value of $\alpha = 10^{-1}$, the focus on timing results in a stable but poor performance level. Above $\alpha = 10^{-1}$, there is a sharp improvement in performance – both the lead time and percentage tardy move to much lower values, which then more or less stabilise above $\alpha = 10^3$. Figure 6(c) presents the standard deviation of lateness results in the pure job shop and general flow shop. The increasing standard deviation at high values of α indicates that some individual orders may become very late when there is an extreme focus on load balancing, i.e. when the timing function is neglected. This effect starts for $\alpha > 10^2$. To take advantage of improved load balancing, our main experiments used $\alpha = 10^3$, where the side effect is still limited. It is remarkable that similar patterns can be observed for the two shop configurations, with changes in performance even occurring at very similar

Figure 4. Simulation results for ILOOR+ vs. all other release methods. (a) Lead time performance in the pure job shop. (b) σ of Lateness performance in the pure job shop. (c) Percentage tardy performance in the pure job shop. (d) Lead time performance in the general flow shop. (e) σ of Lateness performance in the general flow shop. (f) Percentage tardy performance in the general flow shop.

values of α . This suggests that the findings should be relevant to a wide range of shops found in practice.

Overall, Figure 6 highlights the trade-off involved when setting the value of α between delivering the maximum number of orders on time and minimising the number of orders that suffer from extreme lateness. Perhaps counter-intuitively, increasing the

focus on load balancing – and not on timing – actually appears to improve delivery performance in terms of the percentage tardy. But giving extreme emphasis to load balancing only leads to a marginal further improvement in the overall percentage tardy and causes more severe delays for individual jobs. The relative weight given to load balancing may therefore differ from one

Figure 5. Sensitivity analysis for the relative weight of α in the pure job shop and general flow shop. (a) Lead time performance in the pure job shop. (b) σ of Lateness performance in the pure job shop. (c) Percentage tardy performance in the pure job shop. (d) Lead time performance in the general flow shop. (e) σ of Lateness performance in the general flow shop. (f) Percentage tardy performance in the general flow shop.

context to another depending on the performance measures that are most important to a particular shop.

5.2. Analysis of the improvements resulting from intermediate pull release

Using an intermediate pull release mechanism helps to prevent work centre idleness – it can be considered to be a special kind of load balancing that supplements periodic release. To demonstrate this, we recorded the percentage of orders triggered by the intermediate pull release mechanism at direct load norm levels δ_w of 3, 4, 5 and 6 time units for LOOR+ and ILOOR+; and with a workload norm of 4, 5, 6 and 7 time units for LUMS COR. These ranges cover the norm levels at which the methods appear to perform the best.

The results of these experiments are presented in Figure 7, where the dashed lines refer to the general flow shop and the full lines refer to the pure job shop. The figure shows that more orders are released by the intermediate pull release mechanism in the general flow shop than in the pure job shop, which is consistent with the findings for LUMS COR in Thürer et al. (2012). All of the curves in Figure 7 are downward sloping from left to right. In

other words, as the norm level is tightened, more and more orders are released by the intermediate pull release mechanism. This is because low workload norm levels increase the risk of starvation in-between periodic releases.

The most important observation from Figure 7 is that the number of orders released by the intermediate pull release mechanism is always lower for ILOOR+ than for LOOR+ or LUMS COR. This can be attributed to the integer linear programming model incorporated in ILOOR+, which should allow for better load balancing than under LOOR or LUMS COR. This reduces the risk of needless starvation. More generally, this suggests once more that there is high potential for improving the original myopic order release selection procedure that is still incorporated in the periodic release mechanism of most WLC methods, including LUMS COR.

6. Conclusions

LOOR has been an important release method in the WLC literature, and it has been incorporated in many ERP systems (Breithaupt, Land, and Nyhuis 2002). Yet the approach has been largely neglected since its robustness was questioned at the end of the 1990s (e.g. Land and Gaalman 1998; Oosterman, Land,

Figure 6. The impact of the relative weight of α on the performance of ILOOR. (a) Impact on the mean lead time. (b) Impact on the percentage tardy. (c) Impact on the standard deviation of lateness.

Figure 7. The percentage of orders triggered by the intermediate pull release mechanism for LOOR+, ILOOR+, and LUMS COR in the pure job shop and general flow shop.

and Gaalman 2000). This paper has revisited LOOR in the light of the current state of the art. An intermediate pull release mechanism has been added to LOOR to prevent premature idleness; and, an integer linear programming model has been applied to LOOR to overcome the myopic nature of LOOR's order release selection procedure and improve the trade-off between timing and load balancing. When the two refinements are combined, the solution can match or even marginally outperform LUMS COR, the best-performing order release method from contemporary WLC literature.

In answer to our first research question – concerning whether adding an intermediate pull release mechanism can improve LOOR – our results have shown that an intermediate pull release mechanism improves both the performance of LOOR (i.e. LOOR+) and the newly developed ILOOR method (i.e. ILOOR+). Thus, advances in the WLC literature in the context of LUMS COR (Thürer et al. 2012) appear to be transferrable to other WLC order release methods like LOOR. This is an important result for LOOR but also has a more general significance.

In answer to our second research question – concerning the trade-off between load balancing and timing – our results have shown that using an integer linear programming model to enhance the load balancing capabilities of LOOR's order release procedure can indeed improve its performance. This shifts

performance to a new frontier; and it is possible to move along this frontier by adjusting the relative weight given to load balancing and timing. Putting strong emphasis on load balancing minimises the percentage of tardy orders, but there should be some limit on this to avoid the negative side effect otherwise experienced by some orders that would be delivered extremely tardy. The use of integer linear programming however does not lead to substantial performance improvements compared to LUMS COR, which underlines the strength of this heuristic.

Beyond these two research questions, the paper makes a broader contribution as the criticisms levelled at LOOR by Land and Gaalman (1998) and Oosterman, Land, and Gaalman (2000) have now been overcome. First, the authors questioned the effectiveness of LOOR when combined with a due date-oriented dispatching rule. We have shown that the refinements can overcome the poor performance of LOOR with the PST rule, which means that the best results can be realised at lower workload levels. And second, the authors questioned the effectiveness of LOOR in the general flow shop. We have adopted the previously untested adjustment proposed in Breithaupt, Land, and Nyhuis (2002) to look only four process steps ahead in the routing of jobs when making load calculations at release. This adjustment makes LOOR and all of its refined variants robust to changes in shop configuration (from the pure job shop to the general flow shop). Overall, the improvements allow LOOR to be used in a wide range of shops found in practice.

6.1. Managerial implications

The refined versions of LOOR lead to improved performance, addressing the shortcomings previously identified. Yet each refined version has its own strengths and weaknesses, which should be considered when embarking on an implementation:

- *ILOOR* overcomes the myopic nature of LOOR's selection procedure and only releases orders periodically, which means the direct load of work centres does not have to be monitored real-time or continuously to trigger the release of orders from the pool. But the integer linear programming element adds to the complexity of the method. In addition, the performance improvement gained by ILOOR is less than that achieved by LOOR+ or ILOOR+. The improvement does however highlight the importance of reconsidering the

orders selected by the myopic release procedures incorporated in other WLC methods. Sometimes, when loads are not balanced, more appropriate order sets can even be found by a planner without the need for advanced models.

- *LOOR±* adds an intermediate pull release mechanism, which is activated when a work centre is starving, to the periodic release procedure. *LOOR+* leads to a considerable improvement in performance over *LOOR*, with the results approaching those achieved by *LUMS COR*. This means that, for a company that has already implemented *LOOR*, it can improve its performance by adding an intermediate pull release element. Any investment in technology required to support the procedure is likely to be repaid over time as the competitiveness and dependability of the company will increase.
- *ILOOR±* leads to the greatest improvement in the performance of *LOOR*. It has the strengths of 'I' and '+' such as enhanced load balancing and starvation avoidance, which allows it to compete on performance with *LUMS COR*. But the approach also has the drawbacks of 'I' and '+' such as greater complexity and a reliance on technology & information feedback. *LUMS COR* also needs continuous feedback of information from the shop floor, but it is arguably a much simpler method, which should aid its implementation in practice. The fact that *ILOOR+*, which employs integer linear programming, is not able to greatly outperform *LUMS COR* demonstrates the power and utility of *LUMS COR*'s relatively straightforward heuristic. When applying *ILOOR+* (or *ILOOR*), practitioners should bear in mind that a lower percentage of tardy orders can be obtained by focusing strongly on load balancing but that too much focus on load balancing will lead to negative side effects – with a small subset of orders delivered extremely late.

6.2. Future research directions

Finally, the paper has the following implications for future research:

- Improving other existing order release methods through load balancing: The results for *ILOOR* show that putting more emphasis on load balancing can improve performance. This is a starting point for improving other existing WLC release methods, potentially including *LUMS COR*. Moreover, it may be possible to develop other simpler methods to improve load balancing than by applying an integer linear programming model.
- Exploring the prevention of premature idleness using periodic release only: A highly effective order release method should combine control of the direct load of work centres at a low level with the prevention of premature idleness. We have shown that adding an intermediate pull release mechanism is one effective way of preventing premature idleness. However, the low percentage of order releases triggered by this mechanism when incorporated in *ILOOR+*

may suggest that better load balancing to avoid premature idleness may also be possible within a periodic release method. This possibility warrants further study.

- Conducting field research to learn from implementations of refined versions of *LOOR*: Further field research is required to validate the effectiveness of the refined versions of *LOOR* presented in this paper in practice. Field researchers could also aid practitioners by developing a roadmap for the implementation of *LOOR* and by supporting cost-benefit analyses of the various order release mechanism options.

Funding

This work was supported by National Natural Science Foundation of China [grant number 70932004G0209]; Leading Academic Discipline Project of Shanghai Dianji University [grant number 10XKJ01].

Notes on contributors

Haoyun Yan is an associate professor in the Industrial Engineering Department at Shanghai Dianji University. She graduated from Fudan University and received her PhD in 2008, majoring in Management Science and Engineering. Her research includes Production Planning and Control, Make-to-Order Manufacturing and simulation. She has published several papers in domestic journals and International Conferences indexed by EI.

Mark Stevenson is a professor of Operations Management at Lancaster University Management School (LUMS), UK. His main research interest is in production planning & control in low-volume/high-variety manufacturing companies. He has published in a number of international journals, including *Production Planning & Control*, and regularly attends conferences such as the European Operations Management Association (EurOMA) conference.

Linda C Hendry is a professor of Operations Management at Lancaster University Management School (LUMS), UK. Her research interests include: (i) manufacturing strategy, planning and control for product customisation contexts; (ii) process improvement approaches, such as Six Sigma and; (iii) global supply chain management, including sustainable sourcing. Linda is a member of the European Operations Management Association (and was on the Board as a member of the Finance Team from 2011 to 2014) and a member of The Institute of Operations Management. She has published extensively in a wide variety of journals, including those that focus on Operations Management, Production and Operational Research.

Martin J Land is an associate professor in the Operations Department at the University of Groningen. He graduated in Econometrics and Operations Research and received a PhD in Management Sciences in 2004. His research interests are in workload control, job shop manufacturing, card-based control systems and production planning and control in general. He is involved in several industrial projects and has published articles on these subjects in national and international journals.

ORCID

Martin J. Land <http://orcid.org/0000-0003-1704-9874>

References

- Bechte, W. 1988. "Theory and Practise of Load-oriented Manufacturing Control." *International Journal of Production Research* 26 (3): 375–395.
- Bechte, W. 1994. "Load-oriented Manufacturing Control Just-in-Time Production for Job Shops." *Production Planning & Control* 5 (3): 292–307.
- Bergamaschi, D., R. Cigolini, M. Perona, and A. Portioli. 1997. "Order Review and Release Strategies in a Job Shop Environment: A Review and a Classification." *International Journal of Production Research* 35 (2): 399–420.
- Bertrand, J. W. M., and J. C. Wortmann. 1981. *Production Control and Information Systems for Component Manufacturing Shops*. Amsterdam: Elsevier Scientific Publishing Company.
- Breithaupt, J. W., M. J. Land, and P. Nyhuis. 2002. "The Workload Control Concept: Theory and Practical Extensions of Load Oriented Order Release." *Production Planning and Control* 13 (7): 625–638.
- Cigolini, R., and A. Portioli-Staudacher. 2002. "An Experimental Investigation on Workload Limiting Methods within ORR Policies in a Job Shop Environment." *Production Planning & Control* 13 (7): 602–613.
- Fernandes, N. O., M. J. Land, and S. Carmo-Silva. 2014. "Workload Control in Unbalanced Job Shops." *International Journal of Production Research* 52 (3): 679–690.
- Irastorza, J. C., and R. H. Deane. 1974. "A Loading and Balancing Methodology for Job Shop Control." *AIIE Transactions* 6 (4): 302–307.
- Land, M. J. 2004. "Workload Control in Job Shops: Grasping the Tap." PhD Thesis, University of Groningen. Ridderkerk: Labyrint Publications.
- Land, M. J. 2006. "Parameters and Sensitivity in Workload Control." *International Journal of Production Economics* 104 (2): 625–638.
- Land, M. J., and G. J. C. Gaalman. 1996. "Workload Control Concepts in Job Shops: A Critical Assessment." *International Journal of Production Economics* 46–47: 535–548.
- Land, M. J., and G. J. C. Gaalman. 1998. "The Performance of Workload Control Concepts in Job Shops: Improving the Release Method." *International Journal of Production Economics* 56–57: 347–364.
- Lödding, H. 2013. *Handbook of Manufacturing Control*. Translated by Rett Rossi. Berlin:Springer.
- Melnyk, S. A., and G. L. Ragatz. 1989. "Order Review/Release: Research Issues and Perspectives." *International Journal of Production Research* 27 (7): 1081–1096.
- Oosterman, B., M. J. Land, and G. J. C. Gaalman. 2000. "The Influence of Shop Characteristics on Workload Control." *International Journal of Production Economics* 68 (1): 107–119.
- Shafaei, R., and P. Brunn. 2000. "Workshop Scheduling Using Practical (Inaccurate) Data Part 3: A Framework to Integrate Job Releasing, Routing and Scheduling Functions to Create a Robust Predictive Schedule." *International Journal of Production Research* 38 (1): 85–99.
- Soepenbergh, G. D., M. J. Land, and G. J. C. Gaalman. 2012. "Adapting Workload Control for Job Shops with High Routing Complexity." *International Journal of Production Economics* 140 (2): 681–690.
- Stevenson, M., L. C. Hendry, and B. G. Kingsman. 2005. "A Review of Production Planning and Control: The Applicability of Key Concepts to the Make-to-Order Industry." *International Journal of Production Research* 43 (5): 869–898.
- Tatsiopoulos, I. P., and B. G. Kingsman. 1983. "Lead Time Management." *European Journal of Operational Research* 14 (4): 351–358.
- Thürer, M., M. Stevenson, and C. Silva. 2011. "Three Decades of Workload Control Research: A Systematic Review of the Literature." *International Journal of Production Research* 49 (23): 6905–6935.
- Thürer, M., M. Stevenson, C. Silva, M. J. Land, and L. Fredendall. 2012. "Workload Control (WLC) and Order Release: A Lean Solution for Make-to-Order Companies." *Production & Operations Management* 21 (5): 939–953.
- Thürer, M., M. Stevenson, C. Silva, M. J. Land, L. Fredendall, and S. A. Melnyk. 2014. "Lean Planning and Control for Make-to-Order Companies: Integrating Customer Enquiry Management and Order Release." *Production & Operations Management* 23 (3): 463–476.
- Wein, L. M. 1990. "Scheduling Networks of Queues: Heavy Traffic Analysis of a Two-Station Network with Controllable Inputs." *Operations Research* 38 (6): 1065–1078.
- Wein, L. M., and P. B. Chevalier. 1992. "A Broad View of the Job-shop Scheduling Problem." *Management Science* 38 (7): 1018–1033.
- Wiendahl, H. P. 1995. *Load Oriented Manufacturing Control*. Berlin Heidelberg: Springer Verlag.
- Wiendahl, H. P., J. Glässner, and D. Petermann. 1992. "Application of Load Oriented Manufacturing Control in Industry." *Production Planning & Control* 3 (2): 118–129.