University of Groningen

Concerning Workload Control and Order Release

Thürer, Matthias; Land, Martin J.; Stevenson, Mark; Fredendall, Lawrence D.; Godinho Filho, Moacir

Published in:
Production and Operations Management

DOI:
[10.1111/poms.12304](https://doi.org/10.1111/poms.12304)

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:
2015

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

Citation for published version (APA):

Thürer, M., Land, M. J., Stevenson, M., Fredendall, L. D., & Godinho Filho, M. (2015). Concerning Workload Control and Order Release: The Pre-Shop Pool Sequencing Decision. *Production and Operations Management*, 24(7), 1179-1192. <https://doi.org/10.1111/poms.12304>

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.

Concerning Workload Control and Order Release: The Pre-Shop Pool Sequencing Decision

Matthias Thürer

Jinan University, Huangpu Road, 601, 510632 Guangzhou, China, matthiasthurer@workloadcontrol.com

Martin J. Land

Department of Operations, Faculty of Economics and Business, University of Groningen, 9700, AV Groningen, The Netherlands
 m.j.land@rug.nl

Mark Stevenson

Department of Management Science, Lancaster University Management School, Lancaster University, LA1 4YX, Lancaster, UK
 m.stevenson@lancaster.ac.uk

Lawrence D. Fredendall

Clemson University, Department of Management, 101 Sirrine Hall, Clemson South Carolina 29634-1305, USA, flawren@clemson.edu

Moacir Godinho Filho

Department of Industrial Engineering, Federal University of São Carlos, Rodovia Washington Luís, km 235,
 Caixa Postal 676 – Monjolinho, 13565905, São Carlos, Brazil, moacir@dep.ufscar.br

Every production planning concept that incorporates controlled order release will initially withhold jobs from the shop floor and create a pre-shop pool. Order release is a key component of the Workload Control concept that aims to maintain work-in-process within limits while ensuring due dates are met. Order release includes two decisions: (i) a sequencing decision that establishes the order in which jobs are considered for release; and, (ii) a selection decision that determines the criteria for choosing jobs for release. While selection has received much research attention, sequencing has been largely neglected. Using simulation, this study uncovers the potential for performance improvement in the sequencing decision and improves our understanding of how order release methods should be designed. Although most prior studies apply time-oriented sequencing rules and load-oriented selection rules, analysis reveals that load balancing considerations should also be incorporated in the sequencing decision. But an exclusive focus on load balancing is shown to increase mean tardiness and, paradoxically, require high workloads. A new sequencing rule is developed that only balances loads when multiple orders become urgent. It avoids high mean tardiness and allows the shop to operate at a low workload level. At the same time, the percentage tardy is reduced by up to 50% compared to a purely time-oriented rule. The findings have implications not only for Workload Control but for any concept that features order release control, such as ConWIP and Drum-Buffer-Rope.

Key words: order release; pre-shop pool sequencing rule; workload control; simulation

History: Received: April 2013; Accepted: September 2014 by Panos Kouvelis, after 1 revision.

1. Introduction

This study examines the performance of order release control—one of the main functions of production planning and control (e.g., Bertrand and Wijngaard 1986, Zäpfel and Missbauer 1993)—in job shop-like production environments typical of many small and medium sized make-to-order companies. When order release control is applied, jobs do not enter the shop floor directly. Instead, they are retained in a pre-shop pool and released using criteria that allow the shop to achieve certain performance targets, for example,

to restrict the level of work-in-process inventory and/or maximize due date adherence. Well-known approaches to order release control include: Kanban (e.g., Shingo 1989); Drum-Buffer-Rope (e.g., Goldratt and Cox 1984); Constant Work-in-Process (ConWIP, e.g., Spearman et al. 1990); and, Workload Control (e.g., Land and Gaalman 1996). Such control systems are widely applied in practice. For example, Wisner (1996) reported that about 40% of US machine shops in their sample used some form of order release control, and White et al. (1999) reported that about 50% of small and 70% of large companies use Kanban

control. This study concentrates on the Workload Control concept that was developed for such high-variety contexts as the make-to-order job shop (e.g., Stevenson et al. 2005, Thürer et al. 2014, Zäpfel and Missbauer 1993).

An effective order release control mechanism for job shops combines two functions that determine performance (e.g., Land 2006): (i) a load balancing function, so workloads are not only kept within limits or norms but are distributed fairly evenly across resources (Thürer et al. 2012); and (ii) a timing function, so jobs are released such that their due dates can be met. We examine how these two functions are affected by the sequence in which jobs are considered for release—an important aspect of order release neglected by prior research. This will have direct implications for the Workload Control concept, but it will also contribute to the design of all other concepts that incorporate order release control yet lack load balancing qualities (Germs and Riezebos 2010).

When designing an order release method for Workload Control, the first decision is to establish when order release should occur. Sabuncuoglu and Karapinar (1999) and Thürer et al. (2012) have demonstrated that continuous release methods (i.e., where release decisions can take place at any moment in time) outperform periodic release methods (i.e., where release decisions are taken at fixed intervals). The second decision concerns the release itself and can be divided into two parts: (i) a sequencing decision that establishes the order in which jobs are considered for release; and, (ii) a selection decision that determines the criteria for choosing a particular job or set of jobs for release from the pool. These two decisions jointly affect the amount of time that a job is released before its due date (the timing function) and the workload balance across work centers (the load balancing function).

Much prior Workload Control order release literature has focused on identifying appropriate criteria for the selection decision that improve workload balancing on the shop floor (e.g., Cigolini et al. 1998, Oosterman et al. 2000, Thürer et al. 2011). More recently, Thürer et al. (2012) combined many small independent improvements from prior research on order release (e.g., from Hendry and Kingsman 1991, Sabuncuoglu and Karapinar 1999, Oosterman et al. 2000) and showed that the resulting release method simultaneously reduces work-in-process and improves tardiness performance. This overcame a previous criticism of the Workload Control literature—that its order release methods reduce work-in-process but only at the expense of deterioration in tardiness performance (e.g., Baker 1984, Kanet 1988, Kim and Bobrowski 1995, Ragatz and Mabert 1988). Thürer

et al.'s (2012) study also improved our understanding of how the selection decision and the decision about when to release orders (continuously vs. periodically) influence order release performance, but it ignored the potential for improvement in the sequencing decision.

In general, the sequencing decision has received very limited research attention. To the best of our knowledge, the only study to date was presented by Fredendall et al. (2010), who found a load-oriented rule (from Philipoom et al. 1993) had a significant positive impact on both shop floor throughput times and lead times (i.e., the throughput time plus the pool waiting time) compared to alternative time-oriented rules. Yet, it is time-oriented sequencing rules that are commonly applied in most theoretical and empirical studies on order release. The implicit assumption of time-oriented rules is that the sequencing decision is only responsible for the timely release of jobs and, hence, only contributes to fulfilling the timing function of order release (Land 2006, Land and Gaalman 1996). Fredendall et al.'s (2010) finding on the performance of load-oriented sequencing challenges this assumption and may point toward a source of further improvement in order release performance that has been thus far overlooked. But Fredendall et al. (2010) performed only a limited comparison of three available rules, and they did not extend their analysis to establish the underlying causes of their performance results to design new rules.

This study builds on the findings of Fredendall et al. (2010) and extends the study by Thürer et al. (2012) to address the lack of research into the effects of sequencing rules on order release performance. It uses simulation to test six sequencing rules—five from the literature and one rule developed in this study—and examines the effect of the pre-shop pool sequencing decision on performance, including the robustness of performance to changes in routing direction, processing time variability, and due date tightness. The order release method presented in Thürer et al. (2012) is used as the basis for the experiments and, unlike in Fredendall et al. (2010), the underlying causes of the performance results are examined in detail. This provides a more general insight of relevance to any concept that incorporates order release control.

The remainder of this study is structured as follows. The order release method and sequencing rules included in this study are introduced in section 2. The simulation model applied to evaluate the performance of the rules is then described in section 3 before the results are presented, discussed and analyzed in section 4. Final conclusions are drawn in section 5, including implications for practice and future research.

2. Literature Review: Order Release and Pre-Shop Pool Sequencing Rules

This research starts with the following question:

What impact does the pre-shop pool sequencing decision have on the performance of Workload Control order release?

Section 2.1 introduces the order release method that will be used in the simulations. Section 2.2 then outlines five sequencing rules currently available in the literature before section 2.3 introduces a new rule that builds on the load-oriented rule presented by Philipoom et al. (1993) and tested in Fredendall et al. (2010).

2.1. Workload Control Order Release

There are many order release methods in the Workload Control literature; for examples, see the reviews by Land and Gaalman (1996), Bergamaschi et al. (1997), and Fredendall et al. (2010). In this study, the LUMS COR (Lancaster University Management School Corrected Order Release) method is used, because it was recently shown to be the best order release solution for Workload Control in practice (Thürer et al. 2012). It is important to note that LUMS COR incorporates both a periodic and a continuous release time element, as described in subsections 2.1.1 and 2.1.2. Periodic release allows the workload to be balanced, whereas continuous release avoids premature work center idleness or starvation (Thürer et al. 2012).

2.1.1. The Periodic Release Time Element of LUMS COR. At fixed (periodic) intervals, a sequencing rule, such as one of those to be described in section 2.2 below, determines a priority value for each job. This priority value dictates the order in which jobs are considered for release from the pool, beginning with the first job in the sequence. The job selection decision proceeds by determining the contribution a job will make to the workload of each work center in its routing.

Early studies on Workload Control (e.g., Bertrand and Wortmann 1981, Hendry and Kingsman 1991) often calculated the (*classical* aggregate) load of a work center as the sum of all the processing times of jobs released but not yet completed by a work center. But this ignored the fact that the amount of work still upstream (indirect load) may vary, depending on the position of a work center in the routing of jobs. Therefore, the load contribution to a work center in LUMS COR is calculated by dividing the processing time of

the operation at a work center by the work center's position in the job's routing. This *corrected* aggregate load method (Oosterman et al. 2000) recognizes that an order's contribution to a work center's direct load is limited to only the proportion of time that an order is at the work center. For example, an order's load contribution at the second work center in its routing is set at 50% of the processing time at this work center; similarly, its load contribution at the third work center is set at 33.33%, and so on. Oosterman et al. (2000) demonstrated that this provides a good estimate of the expected average direct load resulting from a release decision.

Next, the selection decision compares the corrected aggregate workload of each work center against predetermined workload limits or norms. A job is released if the new workload at each work center in the job's routing is below its workload norm; otherwise, the job is retained in the pre-shop pool. The full periodic release procedure can be formulated as follows:

1. A priority value is determined for each job in the set of jobs J in the pool.
2. The job $j \in J$ with the highest priority is considered for release first.
3. If job j 's processing time p_{ij} at the i th operation in its routing—corrected for work center position i —together with the current corrected workload W_s^C at work center s corresponding to operation i fits within the workload norm N_s^C at this work center, that is

$$\frac{p_{ij}}{i} + W_s^C \leq N_s^C \quad \forall i \in R_j,$$

with R_j being the ordered set of operations in the routing of job j , then the job is selected for release, that is removed from J , and its load contribution is included, that is

$$W_s^C := W_s^C + \frac{p_{ij}}{i} \quad \forall i \in R_j.$$

Otherwise, the job remains in the pool and its processing time does not contribute to the work center load.

4. If the set of jobs J in the pool contains any jobs that have not yet been considered for release, then return to Step 2 and consider the job with the next highest priority. Otherwise, the release procedure is complete and the selected jobs are released to the shop.

2.1.2. The Continuous Release Time Element of LUMS COR. In addition to the above periodic release mechanism, LUMS COR also incorporates a continuous workload trigger. If the load of any work

center falls to zero, the first job in the pool sequence with that work center as the first in its routing is released from the pre-shop pool irrespective of whether its release would exceed the workload norms of any work center in its routing. This avoids premature work center idleness or starvation. The workload contribution to a work center affected by a continuous release is calculated using the same corrected aggregate load approach as described above for the periodic element of LUMS COR.

2.2. Pre-Shop Pool Sequencing Rules from the Literature

This study considers two sets of sequencing rules: (i) time-oriented sequencing rules; and (ii) load-oriented sequencing rules. Four time-oriented sequencing rules are examined:

- *First-Come-First-Served (FCFS)*, which sequences jobs according to their time of arrival in the pool. This rule was used, for example, by Park and Salegna (1995), Cigolini et al. (1998), Sabuncuoglu and Karapinar (1999), and Fredendall et al. (2010).
- *Earliest Due Date (EDD)*, which sequences jobs according to their due date. This rule was used, for example, by Ragatz and Mabert (1988), Melnyk and Ragatz (1989), Philipoom and Fry (1999), and Philipoom and Steele (2011).
- *Planned Release Date (PRD)*, which sequences jobs according to planned release dates given by Equation (1) below. Two variants of this rule are used in the literature, where either waiting times or operation throughput times are treated as a constant. This rule was used, for example, by Bechte (1988), Land and Gaalman (1998), Fredendall et al. (2010), and Thürer et al. (2011 and 2012).

$$\tau_j = \delta_j - \sum_{i \in R_j} (a_i + p_{ij}) \text{ or } \tau_j = \delta_j - \sum_{i \in R_j} b_i. \quad (1)$$

τ_j = planned release date of job j ; δ_j = due date of job j ; a_i = constant for estimated waiting time at the i th operation in the routing of a job; b_i = constant for estimated throughput time at the i th operation in the routing of a job.

- *Critical Ratio, based on Total Work Content (CR-TWK)*, which sequences jobs using a ratio equal to the time remaining until the job's due date divided by the total work content (or total processing time) of the job. The lower the critical ratio, the higher the priority. This rule was used by Enns (1995).

While most attention has been on time-oriented sequencing rules, the study by Philipoom et al. (1993)

represents an exception where a load-oriented rule known as the Capacity Slack (CS) rule was presented. This rule—identified as the best-performing sequencing rule in Fredendall et al.'s (2010) comparative study—sequences jobs according to a CS ratio given by Equation (2). The lower the CS ratio of job j (S_j), the higher the priority of job j .

$$S_j = \frac{\sum_{i \in R_j} \frac{p_{ij}}{N_s^A - W_s^A}}{n_j}. \quad (2)$$

The rule integrates three elements into one priority measure: the *workload contribution* of the job (i.e., the processing time of job j at operation i : p_{ij}); the *load gap*, based on an uncorrected workload measure referred to as the classical aggregate load in the Workload Control literature (i.e., the difference between the aggregated workload norm and current workload at work center s corresponding to operation i : $N_s^A - W_s^A$); and, the *routing length* (i.e., the number of operations in the routing of job j : n_j), which is used to average the ratio between the load contribution and load gap elements over all operations in the routing of the job. This rule was used, for example, by Malhotra et al. (1994) and Fredendall et al. (2010).

Philipoom et al. (1993) incorporated the classical aggregate load measure in their load-oriented sequencing rule, CS. But as explained in section 2.1.1 above, loads are measured as corrected aggregated loads in our selection method. This means that the load contribution of job j to the load of the work center performing its i th operation is $\frac{p_{ij}}{t_i}$ instead of p_{ij} . As a consequence, loads and norms have different values in our study (corrected rather than classical aggregate load values), with load gaps in the selection method measured as $N_s^C - W_s^C$ instead of $N_s^A - W_s^A$. To ensure consistency between selection and sequencing, we simply replace the aggregate load elements in Philipoom et al.'s (1993) CS rule with their corrected equivalents, which means that the CS ratio logically transforms into Equation (3). For clarity, the resulting rule is referred to as the Capacity Slack CORrected (CSCOR) rule.

$$S_j^C = \frac{\sum_{i \in R_j} \left(\frac{p_{ij}}{t_i} \right) / (N_s^C - W_s^C)}{n_j}. \quad (3)$$

2.3. The Design of a New Pre-Shop Pool Sequencing Rule

Figure 1 illustrates the two functions of order release that influence delivery performance: first, the load balancing function, which reduces the average lateness and shifts the lateness distribution to the left

Figure 1 The Intended Influence of Order Release Control on the Lateness Distribution: Load Balancing Function vs. Timing Function

(gray line); and second, the timing function, which reduces the dispersion of lateness and narrows the distribution (dashed line). The Workload Control literature has typically attributed the load balancing function entirely to the selection decision and the timing function to the sequencing decision (e.g., Land 2006). But in the light of Fredendall et al.'s (2010) findings, it is argued here that a load-oriented sequencing rule should in fact complement the selection rule and support load balancing. Hence, load-oriented sequencing is our starting point for designing a new rule that attempts to further improve order release performance.

2.3.1. Backdrop to Developing a New Rule: Preliminary Analysis of CSCOR. A closer examination of the CS ratio (Equation 3) reveals that jobs with short routings and short processing times are most likely to be found at the beginning of the CSCOR sequence. Meanwhile, jobs with long routings are most likely to be found in the middle of the sequence. This is because their ratio between workload contribution and load gap is averaged across a larger number of operations, which creates a smaller probability of an extremely low or high priority number. Finally, jobs with short routings but long processing times are most likely to be found at the end of the sequence. Thus, CSCOR integrates two approaches that can be used to reduce average lateness in a job shop: improving load balancing and prioritizing small orders (in terms of routing length and processing time).

First, load balancing is realized by giving priority to those jobs that only fill a small part of the existing load gaps. Thus, a large number of jobs can be released and loads can get relatively close to their norm levels. This postpones the starvation of work centers for as long as possible by reducing the variability of queues in front of work centers. Second, prioritizing small orders also helps to postpone starvation since, for example, giving priority to orders with the shortest processing times at a given work center quickly replenishes successive queues.

Processing the smallest orders first also typically involves the least sacrifice to achieve a reduction in the average order throughput time as a large number of small orders can usually be processed quickly at the expense of delaying just a few large orders. Yet a drawback of CSCOR is that it ignores the urgency of jobs, meaning that large orders may have to wait for a long time before being released, resulting in a high mean tardiness. This suggests that performance could be improved if a timing function were integrated into CSCOR to consider the urgency of orders.

2.3.2. MODCS: A New Rule that Integrates a Timing Function into CSCOR. The preliminary analysis above predicts problems with high mean tardiness for CSCOR. A similar problem was observed in the (shop floor) priority dispatching literature for the shortest processing time dispatching rule. In that literature, it was addressed by combining shortest processing time dispatching with operation due date oriented dispatching to create a modified operation due date dispatching rule (e.g., Baker and Kanet 1983). Following this principle, we introduce a Modified Capacity Slack (MODCS) pre-shop pool sequencing rule that combines CSCOR sequencing with sequencing according to PRDs. MODCS can be summarized as follows:

1. Jobs are divided into two classes: urgent jobs, that is, jobs with a planned release date (refer back to Equation (1)) that falls within the next release period or has already passed; and non-urgent jobs. Urgent jobs will always receive priority over non-urgent jobs.
2. Within the class of urgent jobs, jobs are sequenced according to the CSCOR rule.
3. Then, within the class of non-urgent jobs, jobs are sequenced according to the PRD rule.

MODCS combines time-oriented and load-oriented sequencing into one rule. It uses its load-oriented CSCOR element to speed up the process when multiple jobs become urgent. Meanwhile, it uses its time-oriented PRD element to ensure non-urgent jobs are released, so the mix of released jobs can be produced in time, thereby reducing tardiness.

3. Simulation Model

Simulation is used to examine the effect of the six pre-shop pool sequencing rules—the four time-oriented rules (FCFS, EDD, PRD, and CR-TWK), the load-oriented rule (CSCOR), and the new rule that combines time-oriented and load-oriented elements (MODCS)—on the performance of order release. The shop and job characteristics modeled are outlined in section 3.1. The parameters set for the order release

method and sequencing rules are summarized in section 3.2. The shop floor dispatching rule applied for controlling the progress of orders on the shop floor is described in section 3.3. Finally, the experimental design is outlined and the measures used to evaluate performance are presented in section 3.4.

3.1. Overview of Modeled Shop and Job Characteristics

To examine the robustness of results, three environmental factors are considered: routing direction, processing time variability, and due date tightness. A simulation model of a randomly routed job shop or pure job shop (Melnik and Ragatz 1989) and of a general flow shop (Oosterman et al. 2000) were developed in Python[®] using the SimPy[®] module. Both shops contain six work centers, where each work center consists of a single, constant capacity resource. The routing length of jobs varies uniformly from one to six operations. All work centers have an equal probability of being visited, but a job will visit a particular work center a maximum of once in its routing. This procedure creates an undirected routing vector or sequence that is used for the pure job shop. The routing vector is sorted or directed for the general flow shop, which means that work centers tend to be positioned upstream and downstream of each other.

Operation processing times follow a truncated log-normal distribution (Trietsch et al. 2012) with a truncated mean of 1 time unit and a maximum of 4 time units. Three levels of processing time variability are modeled, with a $cv^2 = 0.25, 0.5$, and 1 (after truncation). The level of 1 is equal to the coefficient of variation of an exponential distribution, which is typically considered to represent “high” processing time variability. The “medium” level of 0.5 is equal to the variability of a 2-Erlang distribution, which is applied in many earlier studies, including Thürer et al. (2012). Finally, the level of 0.25 has been chosen for “low” processing time variability. This level is still sufficient to avoid unrealistic, nearly symmetric distributions, as observed for lower cv^2 levels. The inter-arrival time of orders follows an exponential distribution with a mean of 0.648, which—based on the average processing time and number of work centers in the routing of a job—deliberately results in a utilization level of 90%.

Due dates are assumed to be set exogenously by the customer and modeled by adding a random allowance factor to the order entry time. Three levels of due date tightness are modeled by varying the minimum and maximum values of the random allowance factor. More specifically, the due date allowances are selected from a uniform distribution with minimum and maximum values of either 30 & 45 (tight due dates), 40 & 55 (moderate due dates),

or 50 & 65 (loose due dates) time units. The lowest minimum value will be sufficient to cover a minimum shop floor throughput time corresponding to the maximum processing time (4 time units) for the maximum number of possible operations (6) plus an allowance for the waiting or queuing times of 6 time units. This due date setting approach allows for comparison with earlier studies on order release and Workload Control (e.g., Land et al. 2014, Melnyk and Ragatz 1989, Oosterman et al. 2000, Thürer et al. 2012).

3.2. Order Release and the Alternative Pre-Shop Pool Sequencing Rules

As in previous simulation studies on Workload Control (e.g., Land and Gaalman 1998, Thürer et al. 2012), it is assumed that all jobs are accepted, materials are available and all necessary information regarding shop floor routing, processing times, etc., is known. Jobs flow into a pre-shop pool to await release according to LUMS COR. The time interval between releases for the periodic release element of LUMS COR is set to 4 time units and all of the sequencing rules are tested with nine workload norm levels from 4 to 12 time units.

Table 1 summarizes the six sequencing rules outlined in section 2. The parameters for the rules are also specified in Table 1. Preliminary simulation experiments identified these parameters as yielding the best overall performance.

Note that the CS ratio within the CSCOR and MOD-CS rules could become negative due to the continuous starvation trigger that LUMS COR incorporates. A negative ratio would result in the sequencing rule prioritizing a job that contributes to the workload of an already overloaded work center. Therefore, if the workload of a work center is equal to or exceeds the workload norm, that is

$$N_s^C - W_s^C \leq 0,$$

then the job is positioned at the back of the queue by replacing the component $\left(\frac{p_{ij}}{N_s^C - W_s^C}\right)$ related to this work center in the priority value S_j^C by $\left(\frac{p_{ij}}{M}\right)$, where M is a sufficiently large number.

Finally, as a basis for comparison, we also record the results with no order release control, referred to as “immediate release”. Here, jobs are released onto the shop floor upon their arrival and control is only exercised via a shop floor dispatching rule.

3.3. Shop Floor Dispatching Rule

Dispatching on the shop floor follows operation due dates, that is, the job with the earliest operation due date from the set of jobs queuing in front of a work

Table 1 Summary of the Six Sequencing Rules Applied in This Study

Abbreviation	Full name	Brief description	Parameter
FCFS	First-Come-First-Served	The job that arrived in the pool first is considered for release first	None
EDD	Earliest Due Date	The job with the earliest due date is considered for release first	None
PRD	Planned Release Date	The job with the earliest planned release date is considered for release first	$a = 4$
CR-TWK	Critical Ratio-Total Work Content	The job with the lowest critical ratio of slack and total work content (or total processing time) is considered for release first	None
CSCOR	Capacity Slack Corrected	The job with the lowest capacity slack ratio (see Equation 3) based on the corrected workload measure is considered for release first	None
MODCS	Modified Capacity Slack	Jobs are divided into two classes: urgent, that is, jobs with a planned release date that falls into the next release period or has already passed; and non-urgent. Urgent jobs are considered for release first according to CSCOR. Non-urgent jobs are then considered according to PRD	$a = 4$ (low), $a = 5$ (med.), $a = 6$ (high)

center is processed first. When a job is released, operation due dates are determined following Equation (4) below. This procedure is based on Land et al. (2014) and is especially suitable when order release control is applied as it accounts for deviations from the schedule due to order release.

$$\text{if } (\delta_j - t_j^r) \geq 0; \delta_{ji} = t_j^r + i \cdot \frac{(\delta_j - t_j^r)}{n_j} \quad \forall i \in R_j \quad (4)$$

$$\text{else if } (\delta_j - t_j^r) < 0; \delta_{ji} = t_j^r; \quad \forall i \in R_j.$$

δ_{ji} = operation due date of job j at the i th operation in its routing

t_j^r = release date of job j

3.4. Experimental Design and Performance Measures

In summary, the experimental factors are the: six different sequencing rules, nine levels of workload norm for the LUMS COR selection decision (4–12 time units), two shop types or routing characteristics (the pure job shop and the general flow shop, that is, undirected and directed routing), three levels of processing time variability, and three levels of due date tightness. A full factorial design was used with 972 cells, where each cell was replicated 100 times. Results were collected over 10,000 time units following a warm-up period of 3000 time units to avoid start-up effects. These simulation parameters are in line with those used in previous studies that have applied similar job shop models (e.g., Land 2004, Oosterman et al. 2000, Thürer et al. 2012). The parameters have been tested and were found to be sufficient for the experiments with the longest initial transients.

Finally, the principal performance measures considered in this study are as follows: (i) the average *shop floor throughput time* or mean time a job spends on the shop floor (i.e., the mean of the completion date minus the release date); (ii) the *percentage tardy*,

that is, the percentage of jobs completed after the due date, or the proportion of jobs that would be delivered late to the customer; and (iii) the *mean tardiness*, that is, the mean of the tardiness $T_j = \max(0, L_j)$, with L_j being the lateness (i.e., the actual delivery date minus the due date) of job j .

4. Results

Statistical analyses of our results were conducted using an ANOVA based on a block design with the workload norm level as the blocking factor, that is, the nine levels of workload norm were treated as different systems. A block design allowed the main effect of the workload norm level and the interaction effects of the sequencing rules and our three environmental factors to be captured. All main effects and two-way interactions were shown to be statistically significant, as were most three-way interactions. The Scheffé multiple-comparison procedure was used to examine the significance of the differences between the outcomes of the individual sequencing rules. All the sequencing rules except EDD and FCFS were significantly different for at least one performance measure; EDD and FCFS were statistically equivalent on all performance measures (see Table 2).

Section 4.1 provides detailed performance results for the pure job shop under medium processing time variability and moderate due date tightness. This includes an in-depth analysis of the influence of the sequencing rule on the load balancing and timing functions. Section 4.2 then assesses the robustness of the results to our three environmental factors: routing direction, processing time variability, and due date tightness.

4.1. Performance Assessment in the Pure Job Shop

Figure 2a and b presents the percentage tardy and mean tardiness performance vs. the shop floor throughput time results, respectively, for the pure job

Table 2 Results for Scheffé Multiple Comparison Procedure

Sequencing rule (x)	Sequencing rule (y)	Shop floor throughput time		Percentage tardy		Mean tardiness	
		Lower [†]	Upper	Lower	Upper	Lower	Upper
CR-TWK	CSCOR	0.263	0.358	0.029	0.031	−0.389	−0.336
CR-TWK	EDD	−0.050*	0.045	0.000*	0.002	−0.079	−0.026
CR-TWK	FCFS	−0.052*	0.043	−0.001*	0.002	−0.095	−0.042
CR-TWK	MODCS	−0.162	−0.067	0.019	0.022	0.044	0.097
CR-TWK	PRD	−0.198	−0.103	0.001	0.004	−0.043*	0.010
CSCOR	EDD	−0.360	−0.265	−0.030	−0.028	0.283	0.336
CSCOR	FCFS	−0.362	−0.268	−0.031	−0.028	0.267	0.321
CSCOR	MODCS	−0.472	−0.378	−0.011	−0.008	0.406	0.460
CSCOR	PRD	−0.508	−0.414	−0.029	−0.026	0.319	0.373
EDD	FCFS	−0.049*	0.045	−0.002*	0.001	−0.042*	0.011
EDD	MODCS	−0.159	−0.065	0.018	0.021	0.097	0.150
EDD	PRD	−0.196	−0.101	0.000*	0.002	0.010	0.063
FCFS	MODCS	−0.157	−0.063	0.019	0.021	0.112	0.166
FCFS	PRD	−0.193	−0.099	0.000	0.003	0.025	0.079
MODCS	PRD	−0.083*	0.011	−0.019	−0.017	−0.114	−0.060

*Not significant at $\alpha = 0.05$.[†]95% confidence interval.**Figure 2** Shop Floor Throughput Time vs.: (a) Percentage Tardy; and, (b) Mean Tardiness Performance for the Six Sequencing Rules and Immediate Release (IMM) in the Pure Job Shop

shop (i.e., undirected routing) with medium processing time variability and moderate due date tightness. In these performance curves, the left-hand starting point represents the tightest workload norm level (4 time units). The workload norm increases step-wise by moving from left to right in each graph, so that each of the nine data points on a given curve represents one norm level (from 4 to 12 time units). In addition, the results obtained with IMMEDIATE release (IMM) are shown as a single point labeled “X”.

From the figure, it is possible to distinguish between the time-oriented and load-oriented sequencing rules according to their impact on our two measures of tardiness: percentage tardy and mean tardiness. First, in terms of percentage tardy (see Figure 2a), we see that the rules that incorporate a load-oriented element—CSCOR and MODCS—consistently outperform the time-oriented rules. Second, in terms of mean tardiness (see Figure 2b), we see that the pure load-oriented rule, CSCOR, results in a higher mean tardiness than the time-oriented rules. This is because CSCOR does not consider the urgency

of individual jobs. Meanwhile, MODCS—where the urgency of individual jobs is considered—matches, or approaches, the time-oriented rules in terms of mean tardiness. Finally, CSCOR realizes lower throughput times for a specific workload norm level than the other rules; however, MODCS and the time-oriented rules realize their lowest percentage tardy at tighter workload norm levels, that is, lower levels of work-in-process inventory than CSCOR.

In general, the results confirm our preliminary analysis in section 2.3 and highlight the potential of load-oriented sequencing rules to improve performance. For example, our new rule, MODCS, allows the percentage tardy to be reduced by up to 50% compared to PRD—the rule most commonly applied in earlier studies on Workload Control—while simultaneously maintaining similar levels of work-in-process reduction and mean tardiness. How load-oriented sequencing rules achieve improvements in performance will be examined next, before the robustness of the results to our three environmental factors is examined in section 4.2.

4.1.1. Performance Analysis: The Influence of Sequencing on Load Balancing and Timing. To examine the performance impact of sequencing rules on the load balancing and timing functions, Figure 3 presents the distribution of lateness across jobs for CSCOR, MODCS and—as a basis for comparison—PRD, the best-performing time-oriented rule. The lateness results presented are at a workload norm level of 5, 7, and 9 time units (see Figure 3a–c). In addition, the Standard Deviation (SD) of each lateness distribution is given in the legend to Figure 3.

First, Figure 3 confirms that the selection decision has a direct impact on the average lateness, as can be observed from the significant mode shift to the left when the norm is tightened, that is, by moving left from Figure 3c to a. Second, Figure 3 confirms that CSCOR reduces the average lateness more than PRD and MODCS. At a glance, the shape of the curves for CSCOR might suggest that this load-oriented rule also significantly improves the timing function of order release. This, however, is not backed up by the standard deviations of lateness (see legend to Figure 3), which are substantially larger for CSCOR than for PRD or MODCS.

From Figure 2 above, it could be observed that CSCOR achieves its largest reduction in percentage tardy at the relatively high workload norm level of 9 time units. By comparing the lateness distribution of

CSCOR with that for PRD at this workload norm level (Figure 3c), it can be seen that CSCOR improves performance by shifting only the right-hand tail of the distribution. This means it is particularly tardy jobs that are affected by the load balancing element of CSCOR. In contrast, for the lower workload norm level of 5 time units (Figure 3a), the whole distribution is affected.

Figure 4 aids in explaining why it is mainly the relevant tardy jobs that are affected at higher workload norms. The figure gives the frequency distribution of the number of jobs remaining in the pool after the periodic release decision has been made for CSCOR (Figure 4a) and PRD (Figure 4b) at a workload norm of 5, 7, and 9 time units. For the workload norm level of 9 time units, all available jobs could be released from the pool during nearly 50% of the periodic release decisions. It is only in the other periods that the sequencing rule has an influence on which jobs are (or are not) released. The impact of the sequencing rule is therefore limited to the jobs that cannot even be released at a loose workload norm level and it will be these jobs that create the right-hand tail of Figure 3c as they have a high probability of becoming tardy. If the workload norm level is tightened, the CSCOR sequencing rule starts to influence a larger proportion of jobs, shifting a greater part of the distribution to the left (see Figure 3a and b). But this is to

Figure 3 The Lateness Distribution for PRD, CSCOR and MODCS in the Pure Job Shop at a Norm Level of 5, 7 and 9 Time Units

Figure 4 The Frequency of Jobs Remaining in the Pool After Periodic Release at a Norm Level of 5, 7 and 9 Time Units: CSCOR vs. PRD

Figure 5 Shop Floor Throughput Time vs.: (a) Percentage Tardy; and, (b) Mean Tardiness Performance for the Six Sequencing Rules and Immediate Release (IMM) in the General Flow Shop

the detriment of the timing function, that is, some more extreme lateness occurs and the standard deviation of lateness increases. MODCS, on the other hand, clearly avoids this, and the distribution is less affected by the norm level applied. This is because the load balancing element of MODCS is only active when multiple jobs are urgent. Thus, it is mainly jobs that might normally have become tardy that are affected by load balancing. As a consequence, MODCS is able to obtain the same positive effects on percentage tardy performance as CSCOR but at a lower workload norm level.

4.2. Robustness of the Results

The robustness of the results to our three environmental factors—routing direction, processing time variability, and due date tightness—will be examined next in subsections 4.2.1 to 4.2.3, respectively.

4.2.1. The Performance Impact of Routing Direction. Similar conclusions on the relative performance of the sequencing rules to those in the pure job shop can be drawn from our results in the general flow shop, that is, when routings are directed. This is illustrated in Figure 5a and b, which present the percentage tardy and mean tardiness performance vs. the shop floor throughput time for the general flow shop (i.e., directed routing) with medium processing time variability and moderate due date tightness. Only results for one level of processing time variability and due date tightness are shown here as the effects on sequencing rule performance observed were qualitatively similar across these factors.

The main difference that can be seen by comparing Figure 2 from the pure job shop with Figure 5 is a slightly larger reduction in the shop floor throughput time at high norm levels for the general flow shop. However, this difference in performance between the pure job shop and general flow shop is not specific to the sequencing rule and has been identified and explained previously by Thürer et al. (2012).

4.2.2. The Performance Impact of Processing Time Variability. CSCOR reduces the average lateness in a job shop by improving load balancing and prioritizing small orders. Both aspects can take advantage of greater processing time variability. Therefore, compared to the results for medium variability (see Figure 2 above), higher processing time variability causes a stronger relative improvement in the percentage tardy for CSCOR compared to time-oriented rules, even at tight workload norm levels. Meanwhile, the difference in mean tardiness performance becomes less pronounced. This is depicted in Figure 6a and b, which illustrate percentage tardy and mean tardiness performance vs. the shop floor throughput time for high processing time variability. In contrast, when processing time variability is low—see Figure 6c and d—less of a focus on balancing loads is required and CSCOR's percentage tardy performance significantly deteriorates relative to time-oriented rules, particularly if norms are tight. Meanwhile, MODCS is shown to still take advantage of combining timing with load balancing. Only results for the pure job shop with a moderate due date tightness are shown as the effects on sequencing rule performance observed were qualitatively similar across these factors.

4.2.3. The Performance Impact of Due Date Tightness. A similar effect to the one observed above in Figure 6 for processing time variability can be observed for due date tightness. This is illustrated in Figure 7a and b, which summarize the results at tight due dates for the percentage tardy and mean tardiness performance vs. the shop floor throughput time, respectively. However, the norm levels, and consequently the shop floor throughput times, that result in the best performance differ between Figure 6 and Figure 7. This difference can be explained by the different needs resulting from high processing time variability and high due date tightness. While an increase in processing time variability specifically requires a better balancing of loads, tighter due dates simply require the average

Figure 6 Performance Curves for High and Low Processing Time Variability in the Pure Job Shop for the Six Sequencing Rules and Immediate Release (IMM)

Figure 7 Performance Curves for the Six Sequencing Rules and Immediate Release (IMM) at Tight and Loose Due Dates in the Pure Job Shop

lateness to be reduced, with improved balancing being just one way of achieving this.

Figure 8 illustrates the implications for the norm level when average lateness has to be reduced. It gives the average lateness performance for high processing time variability (Figure 8a) and tight due dates (Figure 8b), and shows that the largest reduction in average lateness is always realized at a

relatively tight norm level. As a consequence, the best performance occurs at tighter workload norm levels when due dates are tight. Meanwhile, if due dates are loose, the best tardiness performance is realized at looser workload norms, as illustrated in Figure 7c and d. Again, only results for the pure job shop and medium processing time variability are shown as the interaction effects were very limited.

Figure 8 Average Lateness vs. Shop Floor Throughput Time Performance in the Pure Job Shop with: (a) High Processing Time Variability; and, (b) Tight Due Dates

5. Conclusion

When designing an order release method for Workload Control, the first decision is to establish when jobs are to be considered for release; for example, continuously or periodically. The release decision itself then consists of: (i) a sequencing decision that determines the order in which jobs are considered for release; and, (ii) a selection decision that determines the criteria for choosing a particular job or set of jobs for release from the pool. The decision regarding when jobs are considered for release and the selection decision have received much attention in the literature, culminating in Thürer et al.'s (2012) study that combined several independent improvements from prior research. In contrast, the potential of the sequencing decision for improving order release performance—as recently highlighted by Fredendall et al. (2010)—has been neglected. In prior research, it has been implicitly assumed that the sequencing decision is only responsible for the timely release of jobs and that it does not contribute to load balancing.

Our study questioned this assumption and began by asking: *What impact does the pre-shop pool sequencing decision have on the performance of Workload Control order release?* Our results show that load-oriented rules have the potential to significantly reduce the percentage of tardy jobs compared to time-oriented rules, with a reduction of around 50% observed for our new rule, MODCS, compared to the traditional PRD rule across all experimental settings. Most importantly, we have shown that a sequencing rule is capable of not only contributing to the timing function of order release but also to the load balancing function. This provides important inputs for planning practices and future research on the design of order release mechanisms. The managerial implications are outlined in the next section before future research directions are summarized in section 5.2.

5.1. Managerial Implications

This study has demonstrated that sequencing rules that support the load balancing function of order

release have the potential to realize significant improvements in tardiness performance compared to time-oriented rules. These performance improvements were shown to be robust to a broad spectrum of environmental factors (in terms of routing direction, processing time variability and due date tightness), which should provide the necessary confidence for our findings to be applied in practice.

Improvements in tardiness performance are particularly striking if: (i) processing time variability in a shop is high, as a better balancing of capacity is needed to realize improved performance; and/or (ii) due dates are tight, as this requires a focus on reducing the average lateness to realize improved performance. Under these particular conditions, load balancing considerations become increasingly important. However, our performance analysis also revealed that a focus on load balancing should be limited only to the periods when even loose workload norms would already prevent jobs from being released. As a consequence, a pure load-oriented rule (CSCOR) is dependent on relatively loose workload norms, and consequently long shop floor throughput times, if it is to perform well. At tight norm levels, the focus on load balancing was shown to be triggered needlessly, causing a high mean tardiness and a loss in performance improvement relative to time-oriented rules.

An effective sequencing rule should only activate load balancing when it is necessary, based on increased urgency among jobs. With this in mind, the MODCS rule was developed. It is only when multiple orders are at risk of becoming tardy that MODCS shifts its focus from timing to load balancing, and this allows improvements in tardiness performance to be realized at lower levels of work-in-process on the shop floor compared to CSCOR. As a result, MODCS results in the best performance across all three measures considered in our study: throughput time, percentage tardy, and mean tardiness. This implies that the sequence in which planners and planning tools consider jobs for release should be based on a careful trade-off decision. The decision should focus on load balancing when multiple orders tend to run behind

schedule and on giving priority to the most urgent orders when there is limited pressure on meeting due dates.

5.2. Implications for Future Research

This study has improved our understanding of order release control by examining the underlying mechanisms that determine the effect of the pre-shop pool sequencing rule on performance. However, further research on order release as a whole is needed, considering both when releases take place and the release decision itself (i.e., sequencing and selection) to fully understand how each decision affects performance, how the component parts interact, and how they can be better aligned. Future research should also consider order release rules other than those embedded in the Workload Control concept. In this study, we focused on Workload Control and its order release function as a specific solution for high-variety contexts. But order release is a core function of many production planning and control concepts. Well-known alternative approaches to order release control include ConWIP, Kanban, and Drum-Buffer-Rope. In contrast to Workload Control, these approaches have so far not considered load balancing in their release decision (Germs and Riezebos 2010). Pioneering work is required to explore how to integrate load balancing into these approaches to unlock the potential for performance improvement reported here.

Acknowledgments

We would like to acknowledge FAPESP—*Fundação de Amparo a Pesquisa do Estado São Paulo*—for the financial support provided (2011/10605-1).

References

- Baker, K. R. 1984. The effects of input control in a simple scheduling model. *J. Oper. Manag.* 4(2): 99–112.
- Baker, K. R., J. J. Kanet. 1983. Job shop scheduling with modified due dates. *J. Oper. Manag.* 4(1): 11–22.
- Bechte, W. 1988. Theory and practice of load-oriented manufacturing control. *Int. J. Prod. Res.* 26(3): 375–395.
- Bergamaschi, D., R. Cigolini, M. Perona, A. Portioli. 1997. Order review and release strategies in a job shop environment: A review and a classification. *Int. J. Prod. Res.* 35(2): 399–420.
- Bertrand, J. W. M., J. Wijngaard. 1986. The structuring of production control systems. *Int. J. Oper. Prod. Manag.* 6(2): 5–20.
- Bertrand, J. W. M., J. C. Wortmann. 1981. *Production Control and Information Systems for Component-Manufacturing Shops*. Elsevier Scientific Publishing Company, Amsterdam, The Netherlands.
- Cigolini, R., M. Perona, A. Portioli. 1998. Comparison of order review and release techniques in a dynamic and uncertain job shop environment. *Int. J. Prod. Res.* 36(11): 2931–2951.
- Enns, S. T. 1995. An integrated system for controlling shop loading and work flows. *Int. J. Prod. Res.* 33(10): 2801–2820.
- Fredendall, L. D., D. Ojha, J. W. Patterson. 2010. Concerning the theory of workload control. *Eur. J. Oper. Res.* 201(1): 99–111.
- Germs, R., J. Riezebos. 2010. Workload balancing capability of pull systems in MTO production. *Int. J. Prod. Res.* 48(8): 2345–2360.
- Goldratt, E. M., J. Cox. 1984. *The Goal*. 2nd ed. North River Press Inc., Massachusetts, MA.
- Hendry, L. C., B. G. Kingsman. 1991. A decision support system for job release in make to order companies. *Int. J. Oper. Prod. Manag.* 11: 6–16.
- Kanet, J. J. 1988. Load-limited order release in job shop scheduling systems. *J. Oper. Manag.* 7(3): 44–58.
- Kim, S. J., P. M. Bobrowski. 1995. Evaluating order release mechanisms in a job shop with sequence-dependent set-up times. *Produ. Oper. Manag.* 4(2): 163–180.
- Land, M. J. 2004. *Workload Control in Job Shops: Grasping the Tap*. PhD Thesis, University of Groningen, The Netherlands, Labyrinth Publications, Ridderkerk, The Netherlands.
- Land, M. J. 2006. Parameters and sensitivity in workload control. *Int. J. Prod. Econ.* 104(2): 625–638.
- Land, M. J., G. J. C. Gaalman. 1996. Workload control concepts in job shops: A critical assessment. *Int. J. Prod. Econ.* 46–47: 535–538.
- Land, M. J., G. J. C. Gaalman. 1998. The performance of workload control concepts in job shops: Improving the release method. *Int. J. Prod. Econ.* 56–57: 347–364.
- Land, M. J., M. Stevenson, M. Thürer. 2014. Integrating load-based order release and priority dispatching. *Int. J. Prod. Res.* 52(4): 1059–1073.
- Malhotra, M. K., J. B. Jensen, P. R. Philipoom. 1994. Management of vital customer priorities in job shop manufacturing environments. *Decis. Sci.* 25: 711–736.
- Melnyk, S. A., G. L. Ragatz. 1989. Order review/release: Research issues and perspectives. *Int. J. Prod. Res.* 27(7): 1081–1096.
- Oosterman, B., M. J. Land, G. J. C. Gaalman. 2000. The influence of shop characteristics on workload control. *Int. J. Prod. Econ.* 68(1): 107–119.
- Park, P. S., G. J. Salegna. 1995. Load smoothing with feedback in a bottleneck job shop. *Int. J. Prod. Res.* 33(6): 1549–1568.
- Philipoom, P. R., T. D. Fry. 1999. Order review/release in the absence of adherence to formal scheduling policies. *J. Oper. Manag.* 17(3): 327–342.
- Philipoom, P. R., D. C. Steele. 2011. Shop floor control when tacit worker knowledge is important. *Decis. Sci.* 42(3): 655–688.
- Philipoom, P. R., M. K. Malhotra, J. B. Jensen. 1993. An evaluation of capacity sensitive order review and release procedures in job shops. *Decis. Sci.* 24(6): 1109–1133.
- Ragatz, G. L., V. A. Mabert. 1988. An evaluation of order release mechanisms in a job shop environment. *Decis. Sci.* 19: 167–189.
- Sabuncuoglu, I., H. Y. Karapinar. 1999. Analysis of order review/release problems in production systems. *Int. J. Prod. Econ.* 62: 259–279.
- Shingo, S. 1989. *A Study of the Toyota Production System from an Industrial Engineering Viewpoint*. Productivity Press, Cambridge, MA.
- Spearman, M. L., D. L. Woodruff, W. J. Hopp. 1990. CONWIP: A pull alternative to kanban. *Int. J. Prod. Res.* 28(5): 879–894.
- Stevenson, M., L. C. Hendry, B. G. Kingsman. 2005. A review of production planning and control: The applicability of key concepts to the make to order industry. *Int. J. Prod. Res.* 43(5): 869–898.
- Thürer, M., C. Silva, M. Stevenson. 2011. Optimizing workload norms: The influence of shop floor characteristics on setting workload norms for the workload control concept. *Int. J. Prod. Res.* 49(4): 1151–1171.
- Thürer, M., M. Stevenson, C. Silva, M. J. Land, L. D. Fredendall. 2012. Workload Control (WLC) and order release: A lean

- solution for make-to-order companies. *Prod. Oper. Manag.* **21**(5): 939–953.
- Thürer, M., M. Stevenson, C. Silva, M. J. Land, L. D. Fredendall, S. A. Melnyk. 2014. Lean control for make-to-order companies: Integrating customer enquiry management and order release. *Prod. Oper. Manag.* **23**(3): 463–476.
- Trietsch, D., L. Mazmanyan, L. Gevorgyan, K. R. Baker. 2012. Modeling activity times by the Parkinson distribution with a lognormal core: Theory and validation. *Eur. J. Oper. Res.* **216**(2): 386–396.
- White, R. E., J. N. Pearson, J. R. Wilson. 1999. JIT manufacturing: A survey of implementation in small and large US manufacturers. *Manage. Sci.* **45**(1): 1–15.
- Wisner, J. D. 1996. A study of US machine shops with just-in-time customers. *Int. J. Oper. Prod. Manag.* **16**(7): 62–76.
- Zäpfel, G., H. Missbauer. 1993. New concepts for production planning and control. *Eur. J. Oper. Res.* **67**: 297–320.