



HAL
open science

Dynamical multi-parameter sizing of DDMRP buffers in finite capacity flow-shops

Guillaume Martin, Matthieu Luras, Pierre Baptiste

► **To cite this version:**

Guillaume Martin, Matthieu Luras, Pierre Baptiste. Dynamical multi-parameter sizing of DDMRP buffers in finite capacity flow-shops. *Computers & Industrial Engineering*, 2023, 175, pp.108858. 10.1016/j.cie.2022.108858 . hal-03879499

HAL Id: hal-03879499

<https://imt-mines-albi.hal.science/hal-03879499>

Submitted on 25 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamical multi-parameter sizing of DDMRP buffers in finite capacity flow-shops

Guillaume Martin ^{a,*}, Matthieu Luras ^a, Pierre Baptiste ^b

^a Centre Génie Industriel IMT Mines Albi, Albi, France

^b Polytechnique Montréal, Montréal, Canada

ABSTRACT

The DDMRP (Demand Driven Material Requirements Planning) methodology uses buffer stocks to (i) maintain a high level of service, (ii) stop the spread of uncertainty and (iii) adapt to market changes. According to theory, the size of these buffer stocks should be defined regularly. This sizing involves several parameters and policies to update them, but very little information is available on this subject. We aim to help practitioners choose sizing policies, while maximizing the performance of a given workshop. We have developed an experimental design to compare many combinations of flow-shops and bottleneck constraints, taken from industrial use-cases, and using discrete event simulation. The results show that (i) different degrees of dynamism are needed depending on the performance metric chosen by practitioners, (ii) completely dynamic control does not systematically lead to better performance, and (iii) contrary to what the existing literature on DDMRP suggests, varying buffer sizes may be less effective than fixed ones for an important part of use-cases.

1. Introduction

Over the last decades, supply chain and material management methods faced increasingly complex problems: more volatile markets, increasing uncertainty both from demand and supply, complexity from high degrees of product customization or network of actors (Bennett & Lemoine, 2014). The Demand Driven Material Requirements Planning method (Ptak & Smith, 2016) is a materials management framework that is part of the various recent contributions that were developed to deal with this problem. By building on top of previous successful methods such as MRP, Kanban or Theory of Constraints and adding a few innovations, the DDMRP claims the ability to reach higher performance with simpler management (Miclo, Luras, Fontanili, Lamothe, & Melnyk, 2019).

Using DDMRP requires two main steps: (i) creating a network of buffers to help mitigate uncertainties and (ii) implementing a dynamic adaptation of the sizes of these buffers. Creating the network of buffers (where to place them and what types to use) is the first crucial step of a successful DDMRP implementation, but is outside the scope of this paper. It has been previously addressed in Ptak and Smith (2016) for general principles and in Rim, Jiang, and Lee (2014) with an optimization approach or Pekarčíková, Trebuňa, Kliment, and Trojan (2019) with a consideration for product types. The research work presented in this paper focuses on the second step of the implementation: the sizing of all the physical buffers once placed.

While buffers in DDMRP come in three types (physical, time and capacity), we focus solely on the physical buffers. Dynamic adaptation allows the workshop to follow changes either in the market or in the processes. It is done through frequent update of several parameters which, in turn, determine the buffer sizes (Ptak & Smith, 2016). Our study focuses on this dynamic adaptation mechanism.

Changing the size of a DDMRP buffer is done by adapting several parameters, possibly all at the same time (Ptak & Smith, 2016). These parameters govern the sizes of several zones internal to each buffer that add up to its final size. Updating these parameters is a time-consuming process if done by hand due to potentially computation heavy calculations. Additionally, changing a parameter value upstream in the buffer network may have unforeseen consequences downstream. The original authors of the DDMRP (Ptak & Smith, 2016), as well as some authors in more recent works (Dessevre et al., 2019; Lee & Rim, 2019), have recommended some rules to update these parameters. In this paper, we will refer to these rules as “policies”. As a result, the parameters of a DDMRP buffer can be seen as the output of one or more of these independent policies.

As of now, most sources use a single dynamic parameter to control the sizes over time. Very few authors have investigated the interest of controlling a DDMRP with multiple dynamic parameters, let alone across different workshops configurations. In this context, the research

* Corresponding author.

E-mail address: gmartin@mines-albi.fr (G. Martin).

work presented in this paper aims to answer two questions: (i) what are the main parameters that could be tuned dynamically for DDMRP buffers while maintaining sufficient workshop performance and (ii) what policies or combinations of policies for parameter update give the best performance, and for which variation of flow-shop? We restricted our study perimeter to flow-shop variations as they make up a considerable part of implementations and appear in our industrial use-case.

Section 2 reviews previous work on dynamical parameter setting inside and outside the DDMRP methodology. Section 3 details the characteristics of our use-case and how we designed our experiments. Section 4 summarizes the simulation results statistically and produces a set of recommended policy choices. Lastly, Section 5 details where to go from now with the newly acquired knowledge.

2. Literature review

The literature review is split into two sections. In 2.1, we review the past works on DDMRP that may inform us on how sizing works, and on which parameters or policies influence the sizing. In 2.2, we review the parallels between DDMRP sizing and sizing as defined in other planning methods that preceded DDMRP. Finally, 2.3 concludes on our research questions. These sections should help us identify the research gaps for three questions:

1. what are the parameters influencing the dynamic sizing of a DDMRP physical buffer?
2. which of these parameters are supposed to be dynamically updated according to the DDMRP literature? and,
3. do other fields in literature close to the DDMRP contain elements on dynamical update of these parameters?

2.1. Sizing the DDMRP buffers

Most works on DDMRP use or challenge guidelines given by [Ptak and Smith \(2016\)](#), which serves as a basis for the body of knowledge on the method. They identify a set of equations that, once summed, add up to form the final size of a buffer. Appropriately sized buffers allow to decouple demand from supply, while avoiding having too much in stock. These equations call for four mandatory parameters and one additional parameter:

- the **Average Daily Usage (ADU)**, equivalent to the daily demand of a buffered reference,
- the **Decoupled Lead Time (DLT)**,¹ defined as the longest manufacturing lead time between two consecutive buffers,
- the **Variability Factor (VF)**, a coefficient between 0 and 1 proportional to demand instability,
- the **Lead Time Factor (LTF)**, a coefficient between 0 and 1 inversely proportional to the buffered reference's DLT,
- finally, the **Minimum Order Quantity (MOQ)**, which is an optional parameter accounting for the minimum lot size.

As recently emphasized by [Azzamouri, Baptiste, Dessevre, and Pellerin \(2021\)](#), the founders of the DDMRP gave only a few details on parameter definition. In [Ptak and Smith \(2016\)](#), they set the parameters as follows:

- they recommend using moving averages over the past demands to compute the ADU of buffered references. The value of the average window is left to the end user however. [Ptak and Smith \(2018\)](#) gave an update and proposed to use forecasts in conjunction with past demands in determining the ADU. ADU is the only parameter that is explicitly said to vary over time.

- they set the DLT directly from the reference data-sheet, instead of computing it based on historical data.
- they recommend values for VF and LTF according to a three level scale.
- if possible, they recommend against using MOQ as performance of DDMRP may be hindered by fixed lot sizes.

As a result, in [Ptak and Smith \(2016\)](#), only the ADU is explicitly said to be dynamically controlled. Most of the recent contributions specific to the DDMRP follow these recommendations: [Kortabarria, Apaolaza, Lizarralde, and Amorrortu \(2018\)](#), [Miclo et al. \(2019\)](#) or [Velasco Acosta, Mascle, and Baptiste \(2019\)](#), for example, all study a DDMRP instance under an evolving ADU only. For the remainder of this paper, we will refer to any rule or function to dynamically set a parameter value as a "policy".

Definition 2.1 (Policy). A policy is any rule or function that returns the new value for a dynamically set parameter.

Some authors have tested alternatives parameter setting policies. [Dessevre et al. \(2019\)](#) is the only known source using a variable DLT in a DDMRP context. They evaluate the impact of using different policies to set the value from the previous registered cumulative lead times. They review the impacts of three policies on the performance of a flow shop: (i) keeping the DLT constant thus following other researchers, (ii) using exponential smoothing with different smoothing factors and (iii) setting the DLT as a percentile of registered cumulative lead times for a specific reference in the flow shop. After submitting the flow shop to a perturbation in demand, they conclude that controlling both parameters dynamically is possible but can lead to spiraling effects when bottleneck load is near capacity. But, contrary to the rules from [Ptak and Smith \(2016\)](#), to limit the complexity of their simulation, they do not change the value of the LTF with the value of the DLT.

[Lee and Rim \(2019\)](#) propose an alternative to using the VF parameter. They define the lower zone of the buffer through coefficients of variation of demand and lead time. They achieve lower inventories than [Ptak and Smith \(2016\)](#) with their policy, with high client satisfaction. They also define the DLT as the "average lead time" without further details on how it was measured.

Through their study of the literature, [Azzamouri et al. \(2021\)](#) showed a tendency to use different types of simulations in order to obtain new scientific results on the DDMRP. The most prominent one is Discrete Event Simulation (DES), that allows to model complex workshops and specific DDMRP behaviors. Recent works can be divided into two categories: those that compare performance of different material management methods with DDMRP on the same workshop and those that compare alternative ways to control the DDMRP, still on the same workshop. [Shofa, Moeis, and Restiana \(2018\)](#) use DES to compare MRP and DDMRP in a problem with long purchasing lead times and shows that DDMRP achieves lower average inventory. [Miclo et al. \(2019\)](#) compare MRP, Kanban and DDMRP and show that DDMRP globally outperforms the other two in an assembly workshop. Lastly, [Thürer, Fernandes, and Stevenson \(2020\)](#) compare different production planning and control methods on a multi-stage assembly use-case. They refine results from [Miclo et al. \(2019\)](#) and give insights to choose between different material management methods. [Dessevre et al. \(2019\)](#) and [Lee and Rim \(2019\)](#) belong to the second category. They both evaluate different policies to control the DDMRP on the same use-case. [Dessevre et al. \(2019\)](#) use a complete DES software, while [Lee and Rim \(2019\)](#) use spreadsheet simulation.

[Table 1](#) compares how far policies are detailed in all the previously cited works. It shows that most of the contributions detail how the ADU parameter is set over time; most of the time however, authors assume or elude the policies to set the rest of the parameters. For the remainder of this paper, we define the term "policy" as any function able to propose a new value for a parameter by taking into account either a past time series of data or the present state of the system. It follows the definition in the works of [Koulouriotis, Xanthopoulos, and Tourassis \(2010\)](#).

¹ Our study will focus on these two bold parameters, as shown by the synthesis of the literature further in the paper.

Table 1

Author vs. Parameter in DDMRP works, sorted by date (oldest).

| Reference | Parameter | | | | |
|------------------------------|-----------|-----|----|-----|-----|
| | ADU | DLT | VF | LTF | MOQ |
| Ptak and Smith (2016) | 1 | p | 1 | 1 | p |
| Ptak and Smith (2018) | n | p | 1 | 1 | p |
| Kortabarría et al. (2018) | 1 | - | - | - | - |
| Shofa et al. (2018) | 1 | - | - | - | - |
| Miclo et al. (2019) | 1 | - | - | - | - |
| Velasco Acosta et al. (2019) | 1 | - | - | - | - |
| Dessevire et al. (2019) | 1 | n | - | - | - |
| Lee and Rim (2019) | 1 | p | - | - | - |
| Thürer et al. (2020) | 1 | - | - | - | - |

- not addressed or supposed to follow Ptak and Smith (2016).

¹ single policy.^p partially addressed.ⁿ several policies.

2.2. Dynamic sizing for related parameters outside the DDMRP

The DDMRP methodology inherits from several materials management theories (Ptak & Smith, 2016), it is then fair to turn to these to search for dynamical update policies which could help practitioners size their DDMRP buffers dynamically. We identified two main candidate fields: literature from MRP and from card and token controlled production, such as kanban or CONWIP. In these, we aim to find parallels to our dynamical buffer sizing problem. We also consider close parameter definitions between fields. This subsection is organized thematically to better see parallels between the different parameter types.

Dolgui and Prodhon (2007) and Yeung, Wong, and Ma (1998) define the list of parameters impacting the performance of MRP workshops as:

- the size of safety stocks,
- the length of the safety and planned lead times,
- the lot sizing policies,
- the horizon of the Master Production Schedule (MPS),
- and the planning horizon of the production orders.

Among these five, three are relatively close to DDMRP parameters:

- the safety stock size parameter can be closely linked to the definition of the lower zone of DDMRP buffers (see Ptak and Smith (2016) or Lee and Rim (2019)),
- the lot sizing policy is related to the MOQ parameter,
- and the definitions of safety and planned lead times are close to the definition of the DLT.

Defining the MOQ is, in itself, optional in a DDMRP environment (Ptak & Smith, 2016). It is also a problem of its own, as shown by the numerous papers on the topic 7288 results on the Web Of Science only with the request (“lot sizing” OR “lotsizing” OR “lot-sizing” OR “order quantit*” OR “batch size” OR “lot size*” on July, 30, 2021). For these two reasons, the dynamic lot sizing problem will not be considered further in the paper. Lastly, the MPS planning horizon is very specific to the MRP environment and has no counterpart in a DDMRP setting for the moment. This leaves us with two topics to explore.

Defining safety stocks. Earlier works on the definition of safety stocks in an MRP environment can be found in Orlicky (1974), where they are sized according the standard deviations of errors between real and forecast demand. Inderfurth (2009) extends these works by including a risk factor, either for unstable demands or yields. The risk factor is established according to the uncertainties and acts as a multiplication factor of the safety stock. Grubbström (1999) proposes to use safety stock for end products only. The size of these safety stocks is defined as the difference between the sum of production orders and the sum of products consumed by the customers. For this reference, optimal is decided for the sizes that assure the best margin in the last period. Caridi

and Cigolini (2001) consider safety stocks in an MRP environment for irregular demands. They recommend to use two types of safety stocks: one for day-to-day variance sized without taking spikes into account, and one for unusual demand spikes. Lastly, Axsäter (2015) gives several approaches using either optimization or probabilistic measures to decide the safety stock size.

Setting lead times. When setting lead times, we emphasize the difference between two concepts. Firstly, practitioners may want to estimate the completion date of each individual order with relative degrees of precision, which is called lead time in most systems and is defined as the time between order creation in the system and order delivery. This must be done for each order and may also be referred to as “lead time quotation” in the literature. This, however, is not what is needed in order to define the DLT. DLT definition calls for the second concept: setting lead time as a management parameter in order to size the system (Hopp & Spearman, 2011). As a result, we limit our literature review to setting lead times as management constants and not producing a lead time quotation.

In MRP environments, this concept has two sides: (i) defining the planned lead times and, if necessary, (ii) defining the safety lead times. Yano (1987) defines optimal planned lead times in a two-stage production environment, with stochastic production and supply times. The works are also extended for the N stage problem under the hypothesis of independence of lead times for the different stages. The article uses minimal storing costs and late orders as optimality criteria. Findings include the fact that lead time variance is different at each stage and that, lower in the chain, lead time variation is a combination of variations from the higher stages. The article recommends using a safety lead time on the stage with the most variance and to define the planned lead time afterwards.

Melnyk and Piper (1981) use the same decomposition of planned lead time and safety lead time. They observe that acceptable values for the safety lead times are between two and three times the variation of lead times. They also raise the alert of increasing work-in-process generated by an overestimated safety lead time. They recommend using such a safety lead time for components lower in the bill of materials, especially if they are shared between numerous references.

Buzacott and Shanthikumar (1994) and Molinder (1997) compare using safety stocks and safety lead times in a simulated workshop. They establish several strategies to choose between any of the two: for Buzacott and Shanthikumar (1994), we can use safety lead times if forecast quality is good enough. For Molinder (1997), if there is high variability on demand and little on lead time, using safety stocks is a better choice. In case of high variability for both, use safety lead times.

Hnaïen, Dolgui, and Louly (2009) or Dolgui, Ammar, Hnaïen, Louly, et al. (2013) propose a series of cases for determining optimal planned lead times in an MRP environment with stochastic lead times. The series progressively details the use case of a production stage by continuously reducing the hypotheses. The only hypothesis kept is the knowledge of the distribution of the lead times and their discrete nature. The search for optimal values of the planned lead times is then obtained through discrete optimization.

Axsäter (2006) approaches the problem of defining planned lead times in a multi-stage workshop. The workshop is decomposed into individual stages and the author applies the solution given by Yano (1987) for the single stage to each of them. Conclusions shift towards the difficulty of finding exact numerical solutions to large problems and the need to use heuristics to solve it.

The works from Altendorfer (2015) deal with defining both lot sizes and planned lead times in a single stage shop. Authors provide a decision framework to set both parameters with a desired service level. We recall the warnings from Altendorfer (2015), also found in Sridharan and Berry (1990), on the nervousness induced by too many production parameters changes, however.

Major works on lead time definition also borrow from queuing theory. Central results can be found in Hopp and Spearman (2011), for

example. Most of them derive from the now famous Little’s law (Little, 1961) applied to production lines:

$$\text{Work in process} = \text{Cycle time} * \text{Throughput}$$

For example, Hopp and Spearman (2011) consider planned lead time as the probability of the actual lead time to be under a certain value. That value can be obtained as a quantile of the lead times distribution. It could also be set using normal distributions or t distributions if appropriate (Hopp & Spearman, 2011). Hopp and Spearman (2011) also recommend aggregating the flows when considering a multi-product system. It is done by expressing throughput in costs of goods sold (euros per day) and work in process in euros for example.

2.3. Literature synthesis

Analysis of the works on the DDMRP shows that authors agree on five parameters to define the buffer size, with varying degrees of importance. The majority of authors evaluate the DDMRP’s performance and, consequently, follow the policies prescribed by Ptak and Smith (2016). Few authors have challenged these policies and successfully proposed variants.

The DDMRP methodology claims better performance by regularly updating the buffer sizes through these parameters. Most authors only assume a varying ADU; some include a varying DLT. It appears that no work has been done on combining more than one varying parameter, set by more than one policy.

In the meantime, literature shows that, outside works specific to the DDMRP, there are numerous policies that could be used to set the DDMRP parameters dynamically. As a whole, previous works also show the interest of using simulated workshops to establish new results.

Motivated by the gaps in the literature, we focus our study on dynamically controlling the Average Daily Usage (ADU) and the Decoupled Lead Time (DLT) of a DDMRP workshop, as they are the two most influential parameters for setting buffers sizes (see next section for detailed formulas). We will carry the experiment on simulated versions of several types of workshops to (i) have more general results and (ii) set the appropriate level of workshop complexity and still be able to analyze simulation outputs. This design of experiment approach supported by simulation is found in previous DDMRP works such as Miclo et al. (2019) (see 3.3 for more details). Our main objective is to identify to what extent can a DDMRP workshop be dynamically controlled by the ADU and DLT parameters, and what situations may be disadvantageous.

3. Dynamically sized DDMRP model

We base our study on the use-case of dental prosthesis plant, which manufactures semi-finished blanks. Section 3.1 details the main model assumptions behind our work and the plant capacities. In this section, we also detail the future block factors of our experiments: the various types of shop structures in between stages and the different bottleneck stresses. These block factors help generalize our study outside the use-case by synthesizing specific characteristics into more universal types of problems. In Section 3.2, we add dynamical control policies to our model and explain all possible policy definitions we will test. Section 3.3 explains our design of experiments and the metrics we used in the study.

3.1. Shop characteristics and assumptions

Our study is based on an industrial use-case with well defined buffer stages. Stages are enclosed between DDMRP buffers with several parallel machines in between. The plant produces 142 finished products overall, all from only 2 types of raw materials. Following Thüerer et al. (2020), we did not build a complete simulation of the 142 individual

Table 2
Work center capacities.

| Work center | 0 | 1 | 2 | 3 | 4 | 5 |
|-------------|---|---|---|---|---|---|
| Capacity | 6 | 3 | 3 | 3 | 3 | 2 |

routings and constructed a smaller aggregated model with same features as the original. We also assumed raw material buffers to be finite and refilled according to stochastic lead times (exponential distribution using the reference’s DLT as parameter). Different shop structures are found in the global production process, as will be detailed further in the paper.

In our use-cases, work centers have stochastic set-up and processing times. We used the industrial data to adapt triangular distributions for each product-machine pair. All work centers have parallel machines, with varying capacities (details in Table 2), with perfect yield. Work centers are opened 24 h a day, 5 days a week. In line with the industrial use-case, we assumed batch moves and sequential batch production at the work centers. All the work centers follow the queue discipline recommended by Ptak and Smith (2016): highest priority is given to the order with the lowest buffer penetration ratio (buffer level over buffer size in percentage).

Demand data were aggregated for 10 final products over 1000 worked days from the industrial dataset. As shown in Fig. 1, demand knows two types: (i) relatively stable as for Product 2 on the left, or (ii) with various degrees of lumpiness as for Product 5 on the right. Such data patterns were already encountered in previous DDMRP works. See Pekarčíková et al. (2019) or Dessevre et al. (2019), for example.

3.1.1. Shop structure variations

Along the complete industrial process, three shop structures can be found at various stages. We isolated each of them into a separate buffer stage, in order to assess policy choices for each structure. These structure are close to the ones found in previous works on lead time evaluation, such as Öztürk, Kayahgil, and Özdemirel (2006). We named the structures: type A, type V and type O.

Type A workshop is shown in Fig. 2 on the left. In a type A workshop, all product routes coincide at the last work center. Note that this is not an assembly work center however. Type V workshop is shown in Fig. 2, at the center. In this type of workshop, all routes start at the first work center and then diverge. Lastly, type O workshop is shown in Fig. 2, on the right. In a type O workshop, all routes start at the first work center and end at the last one. Whatever the structure, routes are assigned a single branch and material flow is not shared between branches.

3.1.2. Bottleneck factors

Much like the works of Thüerer et al. (2020), we include several levels of bottleneck occupation in our workshops. The motivation for this choice is twofold: (i) first, the industrial stages that motivated the study have several bottlenecks either because of high occupation of a limited resource work center or because common work centers are clogged by the sequential batch production discipline of the plant and, (ii) we wish to study the system under increasing load levels and refine our recommendations of the best policies for dynamical parameter update.

We consider three levels of bottleneck occupation, noted Low, Medium and High. Corresponding occupation ratios are given in Table 3. Table 4 also defines which machine is the bottleneck in each workshop structure.

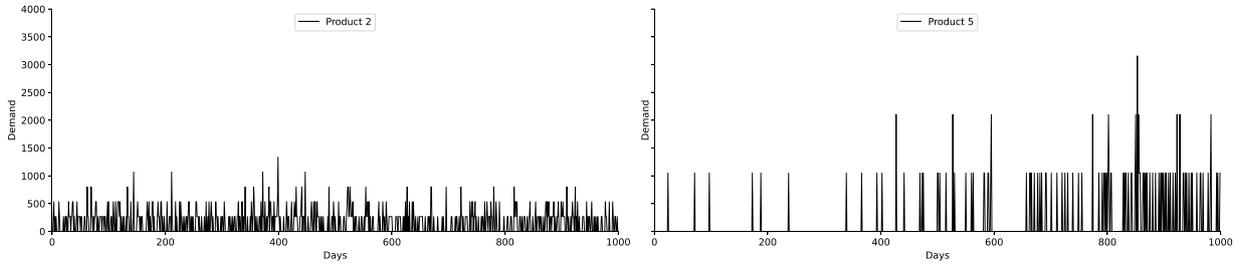


Fig. 1. Two types of demands encountered in the dataset.

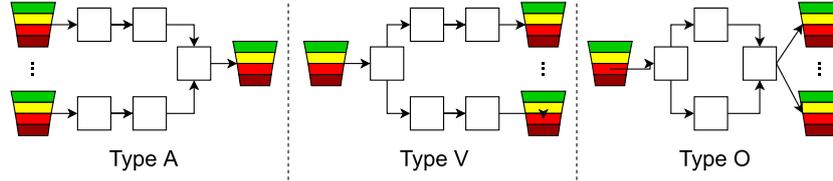


Fig. 2. Types of workshops.

Table 3
Bottleneck occupation ratios.

| Bottleneck level | Low | Medium | High |
|------------------|-----|--------|------|
| Occupation ratio | 50% | 80% | 95% |

Table 4
Workshop VS. bottleneck position.

| Workshop type | Type A | Type V | Type O |
|--------------------|--------|--------|--------|
| Bottleneck machine | Last | First | First |

3.2. ADU and DLT dynamical policies

As explained in 2.3, we wish to investigate a two-parameter dynamical update of a DDMRP workshop. These two parameters relate to: either (i) following the changes in demand with the ADU, or (ii) controlling the internal flow of materials with the DLT. The Eqs. (1) to (4) dictate the standard way of defining zone sizes (Ptak & Smith, 2016) (we use “standard” to remind the reader that the majority authors use these equations, see Lee and Rim (2019) for alternatives). The buffer’s total size is the sum of all zones. For each parameter, we selected three policies, defined hereafter.

$$\text{Green zone} = \text{Max}(\text{ADU} * \text{DLT} * \text{LTF}, \text{MOQ}) \quad (1)$$

$$\text{Yellow zone} = \text{ADU} * \text{DLT} \quad (2)$$

$$\text{Red zone base} = \text{ADU} * \text{DLT} * \text{LTF} \quad (3)$$

$$\text{Red zone safety} = \text{ADU} * \text{DLT} * \text{LTF} * \text{VF} \quad (4)$$

3.2.1. ADU update policies

The first policy is the constant ADU policy: the parameter is set at the beginning of the experiment and never updated. This level serves as a baseline for comparisons. In this setting, ADUs are selected randomly between 50 and 200 units per day. This distribution is based on the average demands from the industrial use-case. For the rest of the paper, this policy is called “Constant ADU” and encoded as “ADU0” in the results analysis.

The second policy is a moving average over past demand. This is the policy recommended by Ptak and Smith (2016). It uses the previous 30 open days to determine the ADU using a moving average function. The 30 days window was determined based on Ptak and Smith (2016): window equivalent to twice the longest lead time in the available data,

15 days in our use-case. We call this policy “Simple moving average” for the rest of the study and encode it as “ADU1”.

The last policy is a moving average including forecasts, as proposed in Ptak and Smith (2018). This policy takes the last 30 open days, adds a 30-day forecast and takes the average of these 60 data points as the new ADU. It is said to better counteract disturbances in demand (Ptak & Smith, 2018). We refer to this policy as “SMAFC” for Simple Moving Average with ForeCasts. We encode it as “ADU2”.

3.2.2. DLT update policies

As in the previous paragraph, the first policy is a constant policy for the DLT. DLTs are selected randomly from the available data-sheets containing 5, 10 or 15 days. This will also serve as a comparison baseline and is called “Constant DLT” and encoded as “DLT0”.

The second policy is extracted from Hopp and Spearman (2011) as there is no mention of a dynamical policy for DLT in Ptak and Smith (2016). For each reference, it returns the 95th percentile of the registered manufacturing times, with a weekly period. This policy is called “Percentile DLT”, and given the code “DLT1”.

The last policy is a modification of the previous one. We call it “Unique DLT” and define it as an extension of the previous policy where all references receive the same DLT value, equal to the 95th percentile of all the registered manufacturing times. Is it proposed in order to evaluate if controlling a stage of the workshop with a single dynamic DLT suffices to ensure good performance. It is also motivated by the fact that slowest references with long lead times may hinder global flow in the workshop (Goldratt & Cox, 1984; Schragenheim & Ronen, 1990). This last policy is encoded as “DLT2”.

3.3. Design of experiments and metrics

In order to evaluate the multiple combinations of policies and workshop layouts, we created a Design Of Experiments (DOE) revolving around Discrete Event Simulation (DES). Several other papers have already showed the efficiency of this approach (see Georgiadis and Politou (2013) and Zhang, Jiang, and Guo (2009), or specifically to DDMRP Miclo et al. (2019) and Thürier et al. (2020)). Our work borrows heavily on the methods of Georgiadis and Politou (2013) which used simulation to help practitioners choose between different parameter policies in a drum–buffer–rope workshop. Using computer simulation is especially beneficial whenever analytical forms of the model are not available or hardly tractable, or when it is necessary to run multiple iterations of a simulation to account for randomness effects. This is exactly our situation as DDMRP methods involve multiple non-linear

Table 5
Experiment setup values.

| Factor type | Factor name | Factor level (Factor Code) | | |
|---------------|-----------------------|----------------------------|------------------------------|-------------------|
| Block factor | Workshop structure | Type A | Type V | Type O |
| | Bottleneck occupation | Low | Medium | High |
| Policy factor | ADU policy | Constant ADU (ADU0) | Simple Moving Average (ADU1) | SMAFC (ADU2) |
| | DLT policy | Constant DLT (DLT0) | Percentile DLT (DLT1) | Unique DLT (DLT2) |

equations and our workshops have stochastic lead times and processing times.

The chosen DOE is a full factorial design, containing two policy factors and two blocking factors. Factors are arranged as follows:

- Blocking factor 1 is the workshop structure,
- Blocking factor 2 is the bottleneck occupation,
- Policy factor 1 is the ADU update policy,
- Policy factor 2 is the DLT update policy.

We summarize the DOE with the experiment setup values in Table 5. We consider three levels for each factor and, as it is a full factorial DOE, we run all possible combinations of these four factors. The factor named “Workshop structure” goes over the three types of workshops described in 3.1.2. The factor named “Bottleneck occupation” takes into account the occupation/capacity ratio in the workshop, also described in 3.1.2. The factor named “ADU Policy” regroups our three ADU control policies. As the names are somewhat long, we gave each level a code, used in the following figures. Lastly, the factor named “DLT Policy” gathers our three DLT control policies. These levels were also given a specific code for easier reading.

Each experiment is evaluated against four metrics adapted from Mico et al. (2019) or Thürer et al. (2020). The first metric is the *average fill rate*, defined as the average percentage of orders directly satisfied from the DDMRP buffers. Fill rate was chosen because our use-case was Make-To-Stock environment.

The next two metrics evaluate inventory performance. The first is *average finished goods inventory* defined as the average number of end items in the physical buffers, thus excluding work-in-process. The second is *average work-in-process inventory* of all buffered references.

The last metric evaluates lead time performance. We consider *average throughput time*, where we collect completion dates of all production orders minus the release date at the first work center.²

Each individual experiment is replicated 100 times with an individual random seed. Each simulation is run for 1000 days and metrics are computed only after a warm-up period of 300 days.

4. Results, recommendations and discussion

We separate the results into two analyses. In Section 4.1, we build trees showing which factors and levels explain the most change in performance for all our metrics. Section 4.2 then details only the best combinations of policies for each of our simulation scenarios. Section 4.3 discusses the interest of this contribution for managing DDMRP workshops and the limits of our results.

4.1. Conditional inference trees of factor importance

For each of the four metrics, we conducted a 4-way analysis of variance, including all interactions. Each ANOVA revealed statistically significant main and interactions effects at a 1% confidence level ($p \leq 0.01$). To further investigate the effects of each factor, we build conditional inference trees (“*ctrees*”) for each metric (Hothorn, Hornik,

² Following the methods from Thürer et al. (2020), we considered *average lead time* too, defined as the difference between order completion and pool entry dates. However results were exactly the same as *average throughput time*. This metric is consequently left out of the study.

Van De Wiel, & Zeileis, 2006). *ctrees* are grown by performing recursive univariate splits of our factors but do so by using the split with the smallest p -value (hence most statistically significant), instead of maximizing the explained variation (such as classic decision trees). Trees have been cropped to a depth of 4 to limit the place of the figures and because further splits were less significant at the 1% level. Each of the subsequent figures contains a *ctree*, the distribution of the output values at the leaves, along with the means. Distributions in a box-plot help the reader evaluate dispersion of the results for the leaves. For level definitions, please refer to Section 3.2. All metrics were rescaled to the [0, 1] scale in order to compare similar ranges, avoiding uselessly high values but keeping the same distributions.

Fig. 3 shows the tree for the Fill Rate. The most important feature is the type of bottleneck occupation in this case. This is revealed by the three zones delimited by:

- left-hand side of the plot and the “Low” bottleneck condition with better Fill Rates between 20% and 90%,
- right-hand side of the plot, further divided into “Medium” (Fill Rate between 13% and 31%) and “High” (between 7% and 17% Fill Rate) bottleneck occupation ratios.

This is no surprise for the Fill Rate, considering the link between sizing the DDMRP physical buffers and queuing theory. This confirms that little can be done in a DDMRP workshop regarding sizing under heavy bottleneck constraints, and that solutions must be searched elsewhere, in capacity increase or load balancing for example. In most branches, the decision on the ADU policy is encountered before the decision on the DLT policy. Across bottleneck levels and workshop types, using variables ADU policies such as ADU1 or ADU2 prove to bring better results for the Fill Rate, without any clear distinction of which method outperforms the other. When in presence of a “Low” bottleneck level, using the DLT2 policy shows better results on average. For “Medium” and “High” levels, the fixed DLT0 policy seems better.

Fig. 4 gives results for the average Finished Goods Inventory. Bottleneck occupation is partly as important as previously, with a “Low” level governing the left-hand part of tree and other levels the right-hand part. However in this case, the choice of DLT policy is more impacting. In all branches, choosing the fixed DLT0 policy ensure lower finished goods inventory, while variable DLT policies result in higher and more varied inventory levels. The choice of ADU policy has less consequences, with a fixed ADU ensuring lower inventory on average. The distribution of workshop types in the tree shows that in some case of type A, inventories are lower. This might be explained by the first work center helping to “starve” the line in this set-up, thus avoiding build-up of material down the line.

Fig. 5 shows the impacts for WIP. This time the choice of DLT policy is the most important: by choosing either a fixed DLT (DLT0) or a unique variable DLT for the stage (DLT2), WIP is greatly reduced, as shown on the left-hand side of the tree. Performance on this side is equivalent for all levels of bottleneck occupation and workshop types. These branches suggest that using DLT0 enables good and stable results, while using DLT2 may allow to reach areas of very good performance if one is willing to account for instability. The variable DLT1 policy gives interesting results in the Low bottleneck and Type V structure, but it is otherwise outperformed, probably because it creates extremely unstable systems.

Effects of factors on average Throughput Time appear in Fig. 6. Here, choosing the fixed DLT0 policy has major consequences: it leads

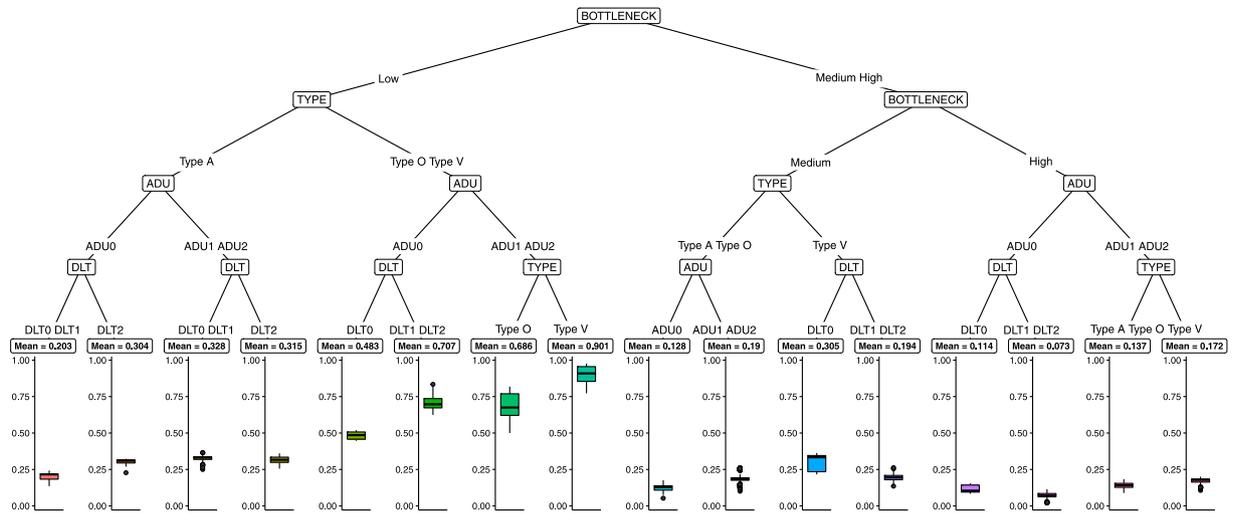


Fig. 3. Conditional tree for factors impacting Fill Rate.

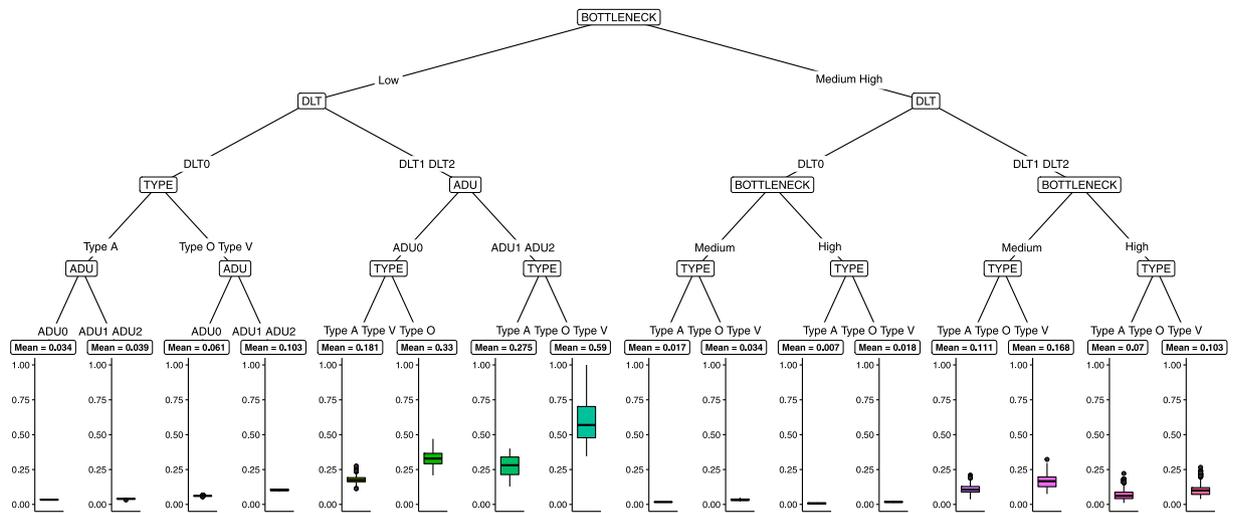


Fig. 4. Conditional tree for factors impacting Finished Goods Inventory.

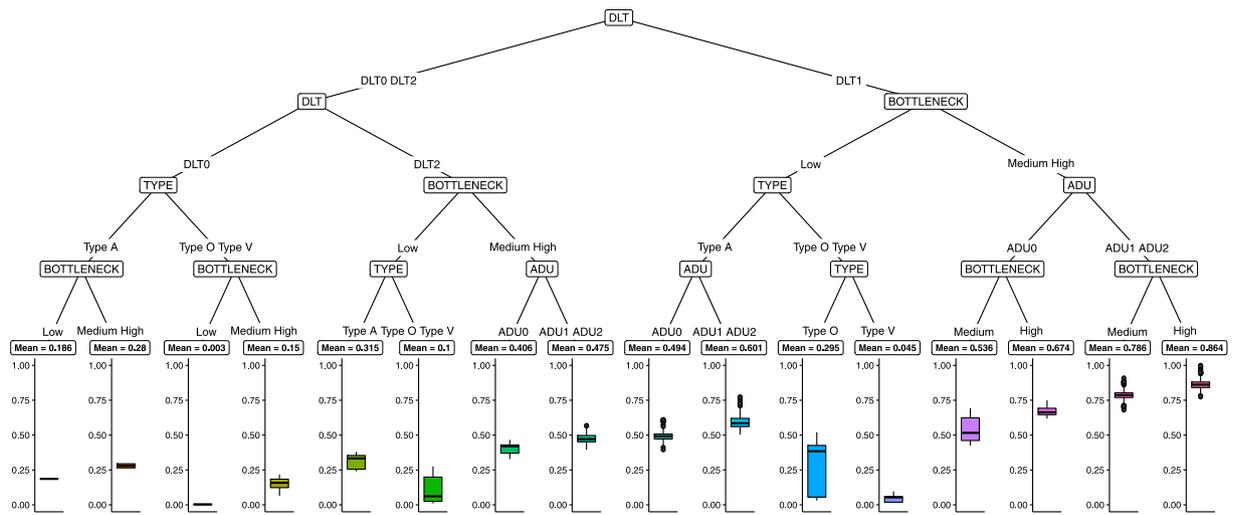


Fig. 5. Conditional tree for factors impacting Work In Process Inventory.

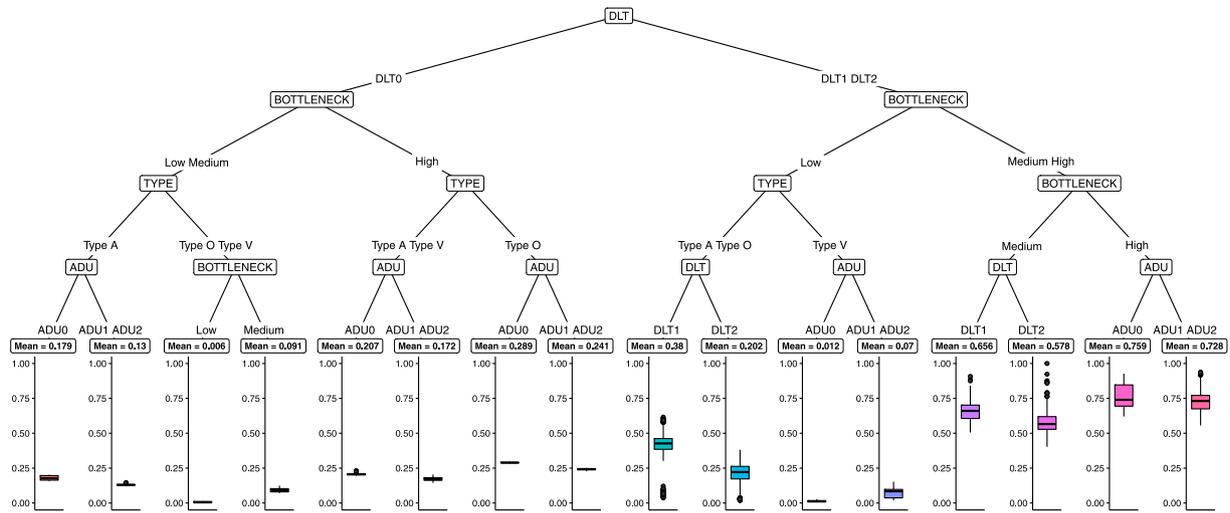


Fig. 6. Conditional tree for factors impacting Throughput Time.

to stabler and lower throughput times across all cases, as shown on the left branches. The exception is the “Low” bottleneck occupation and “Type V” workshop having both ADU and DLT as variable parameters is beneficial to the average throughput time. As in the previous paragraph, we can suspect that a small error in the DLT parameter may have great consequences in the system’s stability whereas, for this metric, the choice of ADU appears much lower in the tree.

4.2. Recommendations on policy choices

We conclude our results with a set of recommendations for practitioners. These recommendations are built using the simulation results. We selected the combination with the best average for each cell. We also confirmed by a one-way ANOVA for each metric that these means were significant at the 1% level. The recommendations are gathered in Table 6. We now go over our six proposed policies and evaluate how they fare performance-wise.

The “Constant ADU” policy ADU0 appears in several cells of Table 6, especially when not considering Fill Rate as the main metric. The “Simple Moving Average” ADU1 and “SMAFC” ADU2 policies appear about as many times but are in majority in the Fill Rate. This may be explained by two aspects of the experiments:

- as each metric is considered independently of others in our study, there is no compromise between all metrics. As a result, a very good combination fill rate-wise may behave poorly on another metric. Consequently, stable combinations of fixed parameter policies may be more efficient than variable ones, especially for finished goods or WIP if the parameters are already close to “optimal”.
- it is possible that the random setting of parameters during initialization was already very favorable. Data came from an industrial dataset where practitioners already have some experience on setting the parameters according to their specific environment.

The overall distribution of the 5 combinations in Table 6 is: (i) 36% for (ADU0, DLT0), (ii) 28% for (ADU1, DLT0), (iii) 25% for (ADU2, DLT0), (iv) 8% for (ADU1, DLT1) and (v) 3% for (ADU2, DLT1). In more than a third of the scenarios, the fully constant solution over 1000 days is the best (but never if Fill Rate is the metric). The combination (ADU1, DLT0), recommended in Ptak and Smith (2016) is the second most frequent and scattered across the four metrics. Fully variable solutions are recommended in only four scenarios (11%). The “Constant DLT” DLT0 clearly outperforms the other and is present in a majority of cells. Very few cells, and only for the Fill Rate metric, call for

Table 6

Policy recommendation based on best simulation results.

| Bottleneck occupation | Workshop type | Fill rate | Finished goods | WIP inventory | Throughput time |
|-----------------------|---------------|------------|----------------|---------------|-----------------|
| Low | Type A | ADU2, DLT0 | ADU0, DLT0 | ADU0, DLT0 | ADU1, DLT0 |
| | Type V | ADU1, DLT1 | ADU0, DLT0 | ADU2, DLT0 | ADU0, DLT0 |
| | Type 0 | ADU2, DLT0 | ADU0, DLT0 | ADU2, DLT0 | ADU0, DLT0 |
| Medium | Type A | ADU1, DLT1 | ADU1, DLT0 | ADU0, DLT0 | ADU1, DLT0 |
| | Type V | ADU2, DLT0 | ADU1, DLT0 | ADU1, DLT0 | ADU0, DLT0 |
| | Type 0 | ADU2, DLT0 | ADU1, DLT0 | ADU1, DLT0 | ADU0, DLT0 |
| High | Type A | ADU1, DLT1 | ADU1, DLT0 | ADU0, DLT0 | ADU1, DLT0 |
| | Type V | ADU2, DLT0 | ADU2, DLT0 | ADU0, DLT0 | ADU1, DLT0 |
| | Type 0 | ADU2, DLT1 | ADU0, DLT0 | ADU0, DLT0 | ADU2, DLT0 |

the “Percentile DLT” DLT1 policy. The “Unique DLT” DLT2 is absent of the recommendations, despite being a major factor in the trees, particularly for average Throughput.

4.3. Discussion

Our main objective is to identify to what extent can a DDMRP workshop be dynamically controlled by the ADU and DLT parameters, and what situations may be disadvantageous.

This study’s objective was to understand to what extent a DDMRP workshop could be dynamically controlled using more parameters than the ADU usually mentioned in the literature. As analytical models were not easily tractable, we had to resort to computer simulation to back up our design of experiments. While, this limits the universal character of our contribution, we introduced blocking factors to take into account multiple workshop types and machine occupation scenarios.

Our results show that, in many cases (36%), stable control is actually a better solution for stable inventory levels and throughput time. This finding goes against the claims of DDMRP being able to adapt to changes and still maintain good performance. Still, in a bit less than a third of the remain cases, the solution proposed by Ptak and Smith (2016) with only a varying ADU is the best one. Contrary to previous claims from Dessevre et al. (2019), dynamically controlling both the ADU and DLT seems to be reserved to a few select cases, and solely when the fill rate is the preferred metric. No clear reason appears for the moment as to why some cases within the fill rate column in Table 6 tolerate multiple dynamic control better than others. Authors from Dessevre et al. (2019) had, however, warned against controlling both parameters in cases where it would introduce instability.

Three aspects then stand out when considering the reach of our findings: the use of independent metrics, the non-sharing of work centers and the trouble of analyzing deep interactions.

We chose to measure performance on four distinct metrics to be able to explain the factor effects independently. This, however, prevented us from having a global (and possibly weighted) view on an aggregated metric. Further studies could be done with the same DOE to refine our analysis. This could prove important as practitioners may look for balanced choices rather than one-sided solutions. This could also help in resolving the discrepancies in policy choice when preferring the fill rate versus the other metrics.

Our use-case required that, outside of the bottleneck work centers, all other work centers were only shared by specific products. This allowed us to define clear branches in the workshop structures. This, however, works against the generalization of our findings; the study could be applied to general job shops, for example.

Lastly, we need to address the analysis complexity limit of our study. The 4-way ANOVA showed meaningful interactions up to the fourth degree, preventing us from drawing conclusions on the main effects of factors. We chose *ctrees* as a way to conduct multiple comparison of all the groups and extract the most important factor choices. Having limited the depth to the fourth level, we lose precision of analysis and the leaves of our trees do not contain the same number of samples. The maximum meaningful depth for the trees according to our data was 7, but it severely hampered reading.

5. Conclusions

We challenged the possibility of controlling two parameters dynamically for several configurations of DDMRP workshops. We tested different combinations of policies to set both the Average Daily Usage (ADU) and Decoupled Lead Time (DLT), while considering different workshop types and bottleneck occupation levels in order to generalize our conclusions at maximum. The variety of our study is mainly motivated by the fact that DDMRP settings come in various shapes and sizes and no two evaluations are really the same. We aimed to start bridging that gap.

Results suggest that two parameter setting (ADU and DLT) may be reserved to very few cases, particularly when considering evaluation metrics independently from each other. This calls for further investigation using a mixture of metrics, either from weighting or from goal programming, for example. Results also hint at the fact that fixed DLT parameters may be sufficient when work centers are not shared. However, in more general workshop structures, this hypothesis may be relaxed, again calling for some more experiments. Lastly, and maybe more surprisingly, we found that a third of cases had better performance when *statically* sized instead of *dynamically*, which goes against common claims of the DDMRP.

We conclude that, for the many scenarios we considered, it may be sufficient to control the sizes of the DDMRP buffers only through the ADU parameter. A good number of cases even suggest keeping it constant, thus simplifying the problem to the accurate sizing of Theory of Constraints or Queuing Theory approaches.

Our study is another step towards understanding the limits and intricacies of applying the DDMRP to real cases. Our findings suggest new avenues for research. Can we derive general recommendations for job shops instead of flow-shops? would composite metrics alter our conclusions? or are there better policies for DDMRP parameter? are some of the still unanswered questions in our field.

CRedit authorship contribution statement

Guillaume Martin: Ideas, Methodology, Software, Validation, Formal analysis, Writing. **Matthieu Laurus:** Ideas, Resources, Data curation, Supervision, Writing – review & editing. **Pierre Baptiste:** Ideas, Supervision, Writing – review & editing.

Data availability

The data that has been used is confidential.

References

- Altendorfer, K. (2015). Influence of lot size and planned lead time on service level and inventory for a single-stage production system with advance demand information and random required lead times. *International Journal of Production Economics*, 170, 478–488.
- Axsäter, S. (2006). Planning order releases for an assembly system with random operation times. In *Stochastic modeling of manufacturing systems* (pp. 333–344). Springer.
- Axsäter, S. (2015). *Inventory control*, Vol. 225. Springer.
- Azzamouri, A., Baptiste, P., Dessevre, G., & Pellerin, R. (2021). Demand driven material requirements planning (DDMRP): A systematic review and classification. *Journal of Industrial Engineering and Management*, 14(3), 439–456.
- Bennett, N., & Lemoine, J. (2014). What VUCA really means for you. *Harvard Business Review*, 92(1/2).
- Buzacott, J. A., & Shanthikumar, J. (1994). Safety stock versus safety time in MRP controlled production systems. *Management Science*, 40(12), 1678–1689.
- Caridi, M., & Cigolini, R. (2001). Buffering against lumpy demand in MRP environments: a theoretical approach and a case study. *Manufacturing Information Systems*, 120.
- Dessevre, G., Martin, G., Baptiste, P., Lamothe, J., Pellerin, R., & Laurus, M. (2019). Decoupled lead time in finite capacity flowshop: a feedback loop approach. In *2019 international conference on industrial engineering and systems management (IESM)* (pp. 1–6). IEEE.
- Dolgui, A., Ammar, O. B., Hnaïen, F., Louly, M.-A., et al. (2013). A state of the art on supply planning and inventory control under lead time uncertainty. *Studies in Informatics and Control*, 22(3), 255–268.
- Dolgui, A., & Prodhon, C. (2007). Supply planning under uncertainties in MRP environments: A state of the art. *Annual Reviews in Control*, 31(2), 269–279.
- Georgiadis, P., & Politou, A. (2013). Dynamic drum-buffer-rope approach for production planning and control in capacitated flow-shop manufacturing systems. *Computers & Industrial Engineering*, 65(4), 689–703.
- Goldratt, E. M., & Cox, J. (1984). *The goal: excellence in manufacturing*. North River Press.
- Grubbström, R. W. (1999). A net present value approach to safety stocks in a multi-level MRP system. *International Journal of Production Economics*, 59(1–3), 361–375.
- Hnaïen, F., Dolgui, A., & Louly, M. A. O. (2009). MRP parameterization under lead times uncertainties: Case of multilevel serial production systems. In *2009 international conference on computers & industrial engineering* (pp. 863–868). IEEE.
- Hopp, W. J., & Spearman, M. L. (2011). *Factory Physics*. Waveland Press.
- Hothorn, T., Hornik, K., Van De Wiel, M. A., & Zeileis, A. (2006). A lego system for conditional inference. *The American Statistician*, 60(3), 257–263.
- Inderfurth, K. (2009). How to protect against demand and yield risks in MRP systems. *International Journal of Production Economics*, 121(2), 474–481.
- Kortabarria, A., Apaolaza, U., Lizarralde, A., & Amorrortu, I. (2018). Material management without forecasting: From MRP to demand driven MRP. *Journal of Industrial Engineering and Management*, 11(4), 632–650.
- Koulouriotis, D., Xanthopoulos, A., & Tourassis, V. (2010). Simulation optimisation of pull control policies for serial manufacturing lines and assembly manufacturing systems using genetic algorithms. *International Journal of Production Research*, 48(10), 2887–2912.
- Lee, C.-J., & Rim, S.-C. (2019). A mathematical safety stock model for DDMRP inventory replenishment. *Mathematical Problems in Engineering*, 2019.
- Little, J. D. (1961). A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9(3), 383–387.
- Melnyk, S. A., & Piper, C. J. (1981). Implementation of material requirements planning: safety lead times. *International Journal of Operations & Production Management*.
- Miclo, R., Laurus, M., Fontanili, F., Lamothe, J., & Melnyk, S. A. (2019). Demand driven MRP: assessment of a new approach to materials management. *International Journal of Production Research*, 57(1), 166–181.
- Molinder, A. (1997). Joint optimization of lot-sizes, safety stocks and safety lead times in an MRP system. *International Journal of Production Research*, 35(4), 983–994.
- Orlicky, J. A. (1974). *Material requirements planning: the new way of life in production and inventory management*. New York, NY, USA: McGraw-Hill, Inc..
- Öztürk, A., Kayalıgil, S., & Özdemirel, N. E. (2006). Manufacturing lead time estimation using data mining. *European Journal of Operational Research*, 173(2), 683–700.
- Pekarcíková, M., Trebuña, P., Kliment, M., & Trojan, J. (2019). Demand driven material requirements planning. Some methodical and practical comments. *Management and Production Engineering Review*, 10, 50–59.
- Ptak, C., & Smith, C. (2016). *Demand driven material requirements planning (DDMRP)*. Industrial Press.
- Ptak, C., & Smith, C. (2018). DDMRP. Demand driven material requirements planning, version 2. Connecticut.
- Rim, S.-C., Jiang, J., & Lee, C. J. (2014). Strategic inventory positioning for MTO manufacturing using ASR lead time. In *Logistics operations, supply chain management and sustainability* (pp. 441–456). Springer.

- Schragenheim, E., & Ronen, B. (1990). Drum-buffer-rope shop floor control. *Production and Inventory Management Journal*, 31(3), 18–22.
- Shofa, M. J., Moeis, A. O., & Restiana, N. (2018). Effective production planning for purchased part under long lead time and uncertain demand: MRP vs demand-driven MRP. In *IOP conference series: materials science and engineering*, Vol. 337. IOP Publishing, Article 012055.
- Sridharan, V., & Berry, W. L. (1990). Master production scheduling make-to-stock products: a framework for analysis. *The International Journal of Production Research*, 28(3), 541–558.
- Thürer, M., Fernandes, N. O., & Stevenson, M. (2020). Production planning and control in multi-stage assembly systems: an assessment of Kanban, MRP, OPT (DBR) and DDMRP by simulation. *International Journal of Production Research*, 1–15.
- Velasco Acosta, A. P., Mascle, C., & Baptiste, P. (2019). Applicability of demand-driven MRP in a complex manufacturing environment. *International Journal of Production Research*, 1–13.
- Yano, C. A. (1987). Setting planned leadtimes in serial production systems with tardiness costs. *Management Science*, 33(1), 95–106.
- Yeung, J., Wong, W., & Ma, L. (1998). Parameters affecting the effectiveness of MRP systems: a review. *International Journal of Production Research*, 36(2), 313–332.
- Zhang, H., Jiang, Z., & Guo, C. (2009). Simulation-based optimization of dispatching rules for semiconductor wafer fabrication system scheduling by the response surface methodology. *International Journal of Advanced Manufacturing Technology*, 41(1–2), 110–121.