

# To Wave Or Not To Wave? Order Release Policies for Warehouses with an Automated Sorter<sup>1</sup>

J eremie Gallien<sup>2</sup> and Th eophane Weber<sup>3</sup>

September 8, 2008

## Abstract

Wave-based release policies are prevalent in warehouses with an automated sorter, and take different forms depending on how much waves overlap and whether the sorter is split for operating purposes. Waveless release is emerging as an alternative policy adopted by an increasing number of firms. While that new policy presents several advantages relative to waves, it also involves the possibility of gridlock at the sorter. In collaboration with a large US online retailer and using an extensive dataset of detailed flow information, we first develop a model with validated predictive accuracy for its warehouses operating under a waveless release policy. We then use that model to compute operational guidelines for dynamically controlling the main parameter of its waveless policy, with the goal of maximizing throughput while keeping the risk of gridlock under a specified threshold. Secondly, we leverage that model and dataset to perform through simulation a performance comparison of wave-based and waveless policies in this context. Our waveless policy yields larger or equal throughput than the best performing wave-based policy with a lower gridlock probability in all scenarios considered. Waveless release policies thus appear to merit very serious consideration by practitioners. Facilities using a non-overlapping wave policy should also consider overlapping waves or a split sorter policy.

## 1 Introduction and Practice Review

Efficiently fulfilling a high volume of small orders chosen from a large number of skus is critical to many online retailers, direct mail-order firms, and retail distributors shipping to many stores on a frequent basis. A common warehouse design in such environments involves an automated sorter allowing different items in the same order to be picked upstream by different workers (Johnson 1998). While automated sorters take different forms, they all typically include an accumulation conveyor and a recirculating loop or "dogtrack" from which items are directed by a diversion mechanism towards individual chutes or lanes (also called "drop points") temporarily assigned to a single order or a small subset of orders with the same shipping destination (Saenz 2002). When all the relevant items have been diverted to a chute, these items can then be packed and/or loaded into a truck as a complete set. Disaggregating a possibly large set of orders into individual item picking instructions distributed

---

<sup>1</sup> The question of this title is from Gilmore (2006b).

<sup>2</sup> Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142.

E-mail: jgallien@mit.edu

<sup>3</sup> Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02142. E-mail: theo\_w@mit.edu

simultaneously in several parts of the warehouse (*batch and zone picking*) and using such a sorter to subsequently re-aggregate these items downstream (since items from the same order may be picked by different workers in different zones) results then in average labor costs per order which can be significantly lower than with other system design alternatives (Choe, Sharp and Serfozo 1992, Petersen 2000, Bragg 2003).

The traditional approach for coordinating the flow of work through the warehouse in such systems is aptly referred to as *wave picking*. In its simplest form, it consists of releasing large batches of orders (the waves) in a sequential manner, so that picking work for a given wave can only start when all the items from the previous wave have been already picked (Choe, Sharp and Serfozo 1992, Petersen 2000). Likewise, items of a given wave are only released into the sorter when all the orders from the previous wave have been already sorted and/or packed (Armstrong, Cook and Saipie 1979, Meller 1997). This approach presents several benefits:

- Using large wave sizes increases the density of items to be picked and thus picking labor productivity (Russell and Meller 2003, Le Duc and de Koster 2007), at least at the beginning and in the middle of the wave (see discussion below);
- Pick lists can be determined for all the workers simultaneously at specific points in time, and can be communicated using simple technology (i.e. paper printout). Indeed, the earliest descriptions of wave picking implementations in industry (e.g. Armstrong, Cook and Saipie 1979) predate by many years the advent of sophisticated computing and wireless communication technologies;
- Blocking effects at the sorter can be completely avoided by ensuring that the number of orders in each wave is less than or equal to the number of sorter lanes (e.g. Armstrong, Cook and Saipie 1979, Johnson and Meller 2002, Owyong and Yih 2006).

However, many sources also discuss several important drawbacks associated with wave picking:

- Because the time for a picker to complete a wave is subject to both predictable and unpredictable variability, even when the work assignment within a wave is balanced across pickers and zones, some idle time may still occur. This may reduce the productivity of picking labor at the end of a wave and ultimately overall picking capacity (e.g. Choe, Sharp and Serfozo 1992, Petersen 2000, see also practitioners' discussions in Gilmore 2006a, Gilmore 2006b and Bradley 2007);

- While large wave sizes typically improve picking labor productivity, they also generate a large buffer of inventory (i.e. cycle stock) between picking and sorting. Such buffer is costly because of the corresponding accumulation conveyor equipment and floor space requirements (Russell and Meller 2003, Bradley 2007). Because of the relationship between wave sizes and number of sorter chutes pointed earlier, wave sizes also drive the sorter purchase cost, which can reach several millions of dollars for large facilities (Johnson and Lofgren 1994, Hinojosa 1996);
- As with any batch processing, waves add a cycle time component to order completion times, which can be an issue in environments with time-sensitive customers such as online retailing (Chew and Tang 1999, Johnson and Meller 2002, Le Duc and de Koster 2007);
- The sequential release into the sorter of non-overlapping waves with a number of orders roughly equal to the number of lanes may result in low capacity utilization. This is because the lanes corresponding to completed and/or packed orders cannot be re-assigned until the end of the current wave (Johnson and Lofgren 1994 and Johnson 1998). This issue is particularly critical during peak periods because sorters are very capital-intensive and thus often constitute the throughput bottleneck (Apple 2006, Perkins 2008, Gilmore 2006a and Bradley 2007);
- Finally, the workload of packers can be fairly concentrated at the end of each wave. This phenomenon can be modeled by using the realistic assumption that items are uniformly distributed within each wave (Hinojosa 1996, Johnson 1998) and considering for each order the expected maximum arrival time into a chute over all its items<sup>4</sup>; given automated sorter are typically not used for single item orders, most chutes only become ready to be packed in the last third of the wave (Hinojosa 1996, 2006), which results in relatively low packing labor productivity at the beginning.

As a result, more sophisticated forms of wave picking have been developed in order to mitigate these problems:

- To reduce pickers' idling at the end of waves, some companies allow different waves to overlap in the picking area, either across zones (Armstrong, Cook and Saïpe 1979) or even within each zone (Owyong and Yih 2006). However, this practice creates the need for a pre-sorting operation downstream in order to separate items from different waves before

---

<sup>4</sup> The last item of an order with  $m$  items is then located on average at a relative position of  $m/(m+1)$  within the wave.

release into the sorter and/or additional buffer space. In addition, differences of wave completion times across pickers and/or zones may accumulate over the course of a shift, unless an effective dynamic re-allocation of labor is performed;

- To increase sorter and packing labor utilization as well as overall throughput, different waves are sometimes also allowed to overlap in the sorter. A first strategy consists of starting the release of each wave as soon as the previous one has reached a specified intermediate completion threshold such as 90% (Bozer, Quiroz and Sharp 1988) or 50% (Johnson and Lofgren 1994). The best threshold to employ is typically determined empirically or through simulation, and seems to vary widely across facilities. As pointed out in Johnson (1998), overlapping waves in the sorter also present some control challenges, in part because of the associated potential blocking effects when all the lanes and/or the recirculating buffer become full (see discussion below). A second strategy consists of splitting the sorter lanes in two, where each half is dedicated to a wave so that packers can work on a completed wave in one half while the next wave is being sorted into the other half (Ruben and Jacobs 1992, Russell and Meller 2003, Perkins 2008). In the ideal operating regime, the sorter completes the processing of that second wave shortly before packers complete the first one, so that packers are not starved for work during the transition between consecutive waves, and complete sorter lanes do not idle for long. Approaching this ideal requires a careful balance of wave completion times across packers, and between the packers and the sorter. To that end, some companies define several categories of workers according to their productivity and assign different numbers of orders or lanes to workers according to their category (with the fastest packers receiving the largest number of orders, see Ruben and Jacobs 1992). However, the pervasive unpredictable variability affecting the order packing and sorting times may easily disrupt this balance, thus reducing labor and/or equipment utilization. In addition, that strategy divides the largest possible wave size by two, which may impact picking labor productivity.

Because even these more sophisticated forms of wave picking still create challenges, a growing number of companies that include American Eagle Outfitters and Green Mountain Coffee have been using in their warehouses an alternative work release control policy referred to as *waveless picking* or *continuous flow picking* (Bradley 2007). While the different implementations of this new policy vary in details (see Forger 2005, Hinojosa 2006, Perry 2007,

Trebilcock 2007, McMahon 2008 and Morris 2008), they all involve the same core principle, which is perhaps best explained through a comparison with traditional wave picking. As discussed earlier, wave picking conceptually involves a first queue of incoming customer orders and a second picking queue corresponding to all the orders covered by the current active picking assignments; whenever the second queue becomes empty, it is replenished at once by an entire batch of a given number of orders (the wave), which is transferred then as a whole from the first queue. In contrast, waveless picking involves the continuous transfer of individual orders from the first queue to the second one, based on a priority ranking of incoming customer orders that is typically based on target shipping dates. The second queue (called a revolving batch or a virtual wave) still has a maximum capacity, which is an important control parameter that we will later refer to as the *revolving batch size*; when that maximum buffer size is reached, any new customer order may only enter the picking queue as another one exits, which occurs when the last one of its items is picked. Pick lists for individual pickers are determined and continually updated in real-time from the picking queue, using a partition of the warehouse storage area into continuous picking loops with a fixed travel direction (the zones), and a dynamic partition of each picking loop between all the pickers assigned to that zone. Specifically, every worker's pick list consists at all times of all the items from orders in the picking queue that are located between his last recorded position and that of the next picker down his picking loop, in the order corresponding to the relative positions of these items along the loop. In addition, this method involves a labor balancing mechanism which continuously evaluates for each zone the expected completion time of the current picking queue by all pickers assigned to that zone, and re-assigns individual pickers across zones whenever imbalances between these completion times exceed specified thresholds. Finally, the buffer upstream of the sorter is used to delay within some limits the induction into the sorter of items belonging to orders with other items further upstream in the process (using for example one of the policies described in Johnson 1998), with the goal of reducing the accumulation time of orders once they are assigned to chutes.

Note that the waveless policy just described critically relies on expensive technology and software, specifically dependable real-time two-way wireless digital communications with every active picker in the warehouse (typically provided by portable devices also including a bar-code reader and an LCD screen), and real-time centralized database management. The

motivation for implementing this novel warehouse control policy however is that while it applies the principles of batch and zone picking and may thus achieve relatively high picking labor productivity (with the density of pick assignments determined by the revolving batch size), it also appears to eliminate some of the inefficiencies associated with wave picking. In particular, no picker is ever starved for work at the end of a wave and, relative to facilities allowing picking waves to overlap, there is no need for a pre-sorting operation. In addition, the rate at which orders become available for packing is more steady because it does not increase towards the end of waves, and completed orders in sorter chutes never need to wait before they can be assigned to a packer. Finally, any urgent incoming order can be assigned for picking almost instantly without waiting until the end of the current wave, and the average completion time of all orders is improved by the elimination of the cycle time before picking and the cycle stock between picking and sorting that are introduced by waves. Indeed, several trade journal articles and corporate white papers point out that waveless picking may generate substantial improvements in both labor costs and throughput relative to traditional wave picking (Hinojosa 2006, Apple 2006, Cooke 2007, Perry 2007), and several support this claim with observations from actual implementations (Forger 2005, Bradley 2007, Morris 2008, McMahon 2008). Industry commentators have also described waveless picking as an application to warehouses of lean manufacturing principles, because of the lot size reduction it entails.

An essential caveat however is that waveless picking no longer involves the release into the sorter of separate batches of a fixed number of orders, which assures that its accumulation space is never exceeded. That new policy therefore creates the potential for severe blocking (Bradley 2007). Specifically, when all the chutes in the sorter are tied up (either by incomplete orders or by complete orders waiting for a packer), upstream congestion can start to build up. As a result, the very items needed to complete orders tying up chutes and relieve this congestion may no longer reach the sorter because of... the same congestion. This phenomenon is known as *gridlock* (Johnson and Lofgren 1994), and because the corresponding recovery procedure is typically long and laborious, it can significantly reduce capacity and productivity (Holste 2008). In the words of Sam Sanders, a warehouse consultant quoted in Bradley (2007): "It's like a game of solitaire. If all the slots in the game are full, the game is over, and you lose. If you have 10,000 SKUs and 1,200 drop points, you can have a lot of

SKUs on the sorter with no place to drop into. If you want to work with continuous flow, you have to be cognizant of this."

The work to be described here started from a collaboration with a leading US online retailer (*our industrial partner*), who had switched all of its warehouses with an automated sorter from wave-based to waveless picking before our interaction began. While our partner reports observing then significant increases in throughput, equipment utilization and labor productivity, it did not initially impose formal guidelines for how managers should dynamically adjust the primary control levers of this new order release policy, and indeed experienced gridlock more frequently than desired during the peak demand seasons following that implementation. This situation motivated in part the two research objectives pursued in this paper, using an extensive dataset of detailed flow information obtained from our industrial partner: (Objective 1) *Develop a quantitative model to generate operational control guidelines for waveless picking during peak demand periods, with the goal of maximizing throughput while keeping the likelihood of gridlock sufficiently low*; and (Objective 2) *Leverage this model to conduct a rigorous performance comparison between wave-based and waveless release policies in the context of our industrial partner's warehouses*. After a discussion of the related literature and our contributions in §2, we describe our work on the first objective stated above in §3. We then present a quantitative model describing wave-based picking in the context of our partner's warehouses in §4, and discuss in §5 the simulation experiments we performed in order to achieve our second objective. Concluding remarks are provided in §6, and the Online Appendix to this paper contains supporting material, including auxiliary results and detailed algorithm statements. Mathematical variables in capital letters refer throughout to random quantities, while those in lower case refer to deterministic quantities. Also, notations with an upper (resp. lower) bar refer to the maximum (resp. minimum) value in an index set or interval, variables in bold refer to vectors or control policies, and new terminology being defined appears in italics.

## **2 Related Literature and Paper Contributions**

We limit our discussion to papers specifically motivated by warehouses with an automated sorter, and we refer the reader to the surveys by de Koster, Le-Duc and Jan Roodbergen (2007) and Gu, Goetschalckx and McGinnis (2007) for recent and extensive reviews of the many design and control problems arising with other types of warehouses. Also, literature

that is methodologically related to our analysis is discussed in subsequent sections.

A first set of papers positively answers the question of whether an automated sorter constitutes a justified design choice in some environments. This question is closely related to the choice of an order picking policy, because the need for sortation only arises when pick lists do not preserve order integrity. Using simple queueing models, Choe, Sharp and Serfozo (1992) compare the order cycle times associated with the three strategies of single order picking, batch picking and batch zone (wave) picking with an automated sorter. With more complex simulation models, Petersen (2000) investigates the policy of sequential zone picking in addition to the three others just mentioned, considering not just order cycle time but also labor requirements. Finally, Russell and Meller (2003) develop a simple deterministic cost model to inform the decision of whether manual or automated sorting should be used in a warehouse operating under wave picking. In the specific environment of online retailer Amazon.com, Bragg (2003) also considers similar simple deterministic cost model to assess various warehouse design alternatives.

A second group describes the use of simulation models to determine various dimensioning and control parameters of warehouses with automated sorters, given specific design objectives. Bozer and Sharp (1985) explore the impact on sorter throughput of the number of sorter lanes and their storage capacity as well as the use of recirculation and the concentration of items from the same order within a wave. That study is complemented by Bozer, Quiroz and Sharp (1988) who also investigate the throughput implications of the wave profile (size, distribution of items per order), the lane assignment policy and the degree to which consecutive waves are allowed to overlap. Finally, Johnson and Lofgren (1994) report the successful use of simulation model decomposition when designing a new warehouse with an automated sorter. This decomposition is enabled by sufficient picking capacity together with the wave release strategy, which effectively decouples the areas of picking and sorting. Also relevant to sorter design decisions is Johnson and Meller (2002), which presents an analytical model predicting the throughput of induction stations used in split-case sorting operations, where workers need to manually place incoming items onto moving trays bound to sorter chutes.

A last set explores more specific operational problems motivated by the use of automated sorters. We first note that, although not reviewed here, some of the papers investigating



the problems of storage assignment, zoning, order batching and picker routing within the order picking area considered in isolation (see de Koster, Le-Duc and Jan Roodbergen 2007) are also potentially relevant to the operations we focus on. However, only a couple develop models that capture the impact of these decisions on the sorting process. Among them, Armstrong, Cook and Saïpe (1979) and Le-Duc and de Koster (2005) formulate mixed integer optimization models to compute order batches, taking into account both discrepancies in wave completion times across zones and limitations of sorter capacity. In the context of module-based fulfilment centers, Owyong and Yih (2006) present a heuristic for modifying pick lists to reduce order consolidation time or *chute-dwell time*, which is the time between the arrivals of the first and last item of a customer order in a sorter chute. Finally, Meller (1997) describes integer programming models for the problem of assigning orders to sorter lanes (lane assignment), assuming that the sequence of incoming items is known. Relaxing this last assumption, Johnson (1998) develops an analytical model predicting the impact of various lane assignment strategies on expected total wave sorting time.

In all the literature just discussed, wave picking is the only operating policy considered for warehouses with an automated sorter. To the best of our knowledge, the present paper is the first to present a mathematical model for the problem of order release control in a warehouse operating under waveless picking (see §1 for background), and to address quantitatively the questions of whether and how gridlock may be avoided under such policy. This study also leads us to investigate and characterize the relationship between the release strategy for picking orders and the sorter operation, which has been consistently described as an important object of research (Johnson 1998, Petersen 2000, Owyong and Yih 2006). Finally, the solution we develop for the order release problem mentioned above enables us to present the first meaningful quantitative performance comparison between wave-based and waveless picking in a specific context.

### **3 Waveless Release Model and Analysis**

As part of the waveless picking policy described in §1, the main daily flow control levers available to our partner include the size of the revolving batch, which can be adjusted a few times per hour, as well as the staffing levels for pickers and packers, which can be adjusted a few times per day (we refer the interested reader to the Online Appendix for a detailed description of our partner’s warehouses). Note that staffing decisions for induction stations

(and more generally their capacity) are not explicitly considered in this section. This is because our partner uses induction stations with automated coordinated induction belts which were designed using realistic throughput models of the type described in Johnson and Meller (2002), so that only few operators are required to achieve a high induction capacity. Also, the differences in possible control change frequencies just stated justify our focus on the revolving batch size under the assumption that staffing levels constitute fixed exogenous parameters.

The revolving batch size affects the overall picking rate in two ways. First and foremost, through the resulting density of items to be picked along the picking paths. Secondly, through the starvation of pickers occurring when the revolving batch size is set to a low value<sup>5</sup>. Fortunately, our partner had developed through extensive empirical studies a simple but reliable model that determines the size of the revolving batch required in its environment in order to generate a specified average overall picking rate under a given staffing level for pickers. However, no formal guidelines were available at the outset of our interaction for dynamically changing the target picking rate as a function of observed process conditions (conveyor system congestion, numbers of complete, incomplete and unassigned sorter chutes) and packers' staffing level<sup>6</sup>. This was mostly problematic during peak demand period, when the pressures for high throughput sometimes resulted in gridlock (see §1).

In the remainder of this section, we present a quantitative model and analysis developed to identify flow control strategies for waveless picking that maximize throughput while keeping the probability of gridlock under an acceptably low level during such peak demand periods. We define our predictive model in §3.1, present related approximate dynamics and discuss their validation in §3.2, state the optimization problem we consider in §3.3 and finally discuss an associated solution algorithm in §3.4.

**3.1 Predictive Model** Our model for predicting congestion in our partner's warehouse as a function of the flow control policy employed is a serial queueing network with three stations and features similar to that found in Gallien and Wein (2001). It is defined as follows:

---

<sup>5</sup> Picking rate in this setting can be conceptually modeled by the throughput of a closed queueing network with a single station having as many servers as there are pickers, and where the number of circulating entities, which corresponds to the revolving batch size, affects the service rate.

<sup>6</sup> Instead, managers tended to apply informal guidelines prescribing to try and stabilize the process around target numbers of complete and incomplete sorter chutes. These target numbers were not supported by any analysis, appeared inconsistent across managers and facilities, and adherence to those guidelines was not enforced.

– Each circulating entity in this network represents a customer order, and the arrival rate of these orders to the network over time is the primary control we ultimately seek to optimize (in §3.3). This input is modeled as a Poisson process whose rate at time  $\tau \geq 0$  is denoted  $\lambda(\tau)$ , and the associated release rate control policy will be noted  $\lambda$ . In the real system,  $\lambda(\tau)$  corresponds to the average current rate at which orders are picked. As in Russell and Meller (2003), we thus do not explicitly model the warehouse layout, stowing policies and picker routing policies employed; however we do consider their aggregate impact on the overall rate at which pickers release items onto the conveyor system. Specifically, this impact is captured by the empirical model mentioned above, which links overall average picking rate with the revolving batch size and the pickers’ staffing level, and thus enables a practical implementation of flow control policies characterized by target order picking rates (converting average order pick rate to average item pick rate is straightforward since the average number of items per order  $\mathbb{E}[M]$  is easily determined). This modular approach is more tractable, but also alleviates the need to explicitly model features which tend to be more idiosyncratic such as the layout of the picking area. In other settings, the relationship between batch size, staffing levels and average overall picking rate may also be determined empirically, or through analytical models of the kind developed by Chew and Tang (1999) and Le-Duc and de Koster (2007). We also note that the arrival process in our model is random, which reflects that the actual system is only partially controllable. That is, specifying the revolving batch size for a given staffing level only implements an average picking rate, from which the current instantaneous picking rate may differ – this stems in practice from variability in the actual density of picks along the picking loops, pickers’ individual productivity, the actual number of items per order, etc. The specific structure assumed for that randomness (Poisson) is motivated by both analytical tractability considerations and the intuitive relevance of Palm’s theorem to this setting. Because the maximum picking rate achievable is limited in practice by various factors, we assume an upper bound  $\bar{\lambda}$  for this control. During the peak demand periods that we are most concerned with, the virtual queue of orders placed by customers but not yet released for picking is always sizeable and never empties. As a result, our assumption that the upper bound  $\bar{\lambda}$  is exogenous seems reasonable. Finally, because of practical database synchronization issues, the frequency at which the revolving batch size may be

adjusted is limited. As a result, we assume that the average picking rate  $\lambda(\tau)$  may only be changed in our model at discrete time points separated by a period  $\delta$  (of the order of a few minutes). As a result, the release policies  $\boldsymbol{\lambda}$  considered effectively implement a discrete sequence of controls  $(\lambda_t)_{t \in \mathbb{N}}$ , where each discrete time period  $t \in \mathbb{N}$  corresponds to the continuous time interval  $[t\delta, (t+1)\delta)$ , i.e.  $\lambda(\tau) = \lambda_t$  for  $\tau \in [t\delta, (t+1)\delta)$ .

- The first station of this queueing model has an infinite number of servers, each with identically distributed service times following a distribution noted  $A$  and representing the *time-to-chute*, or time between the first time at which an item belonging to a customer order is released for picking anywhere in the warehouse and the first time at which an item from that order reaches a sorter chute. The process representing the number of orders undergoing service in this first station is denoted  $X(\tau)$ , and provides a partial measure of the conveyor congestion upstream of the sorter. In the following, we will use the notation  $X_t \triangleq X(t\delta)$ .
- The second station has a finite number of servers equal to the number of sorter chutes  $n$ , each with identically distributed service times following a distribution noted  $B$  and representing the *chute-dwell time* of every order, which is the time that each customer order spends in a chute before all its items are complete. The number of orders undergoing service in this second station thus represents the number of incomplete chutes in the sorter at any point in time, and follows a process denoted  $Y(\tau)$ . As before, we define  $Y_t \triangleq Y(t\delta)$ .
- Finally, the third station represents the packing stage. It has a finite number of servers equal to the number  $w$  of packers assigned to the sorter, each with identically distributed service times representing the *pack-to-pack time*  $C$ , or cycle time experienced by a packer for each customer order (e.g. time spent walking to the next chute + time spent packing). The process representing the number of orders in this station (in queue and in service) is denoted  $Z(\tau)$ , which thus corresponds to the number of complete chutes in the sorter at any point in time. Its values at the discrete time points  $(t\delta)_{t \in \mathbb{N}}$  are also denoted  $Z_t \triangleq Z(t\delta)$ .

The considerations that led us to formulate the predictive model just described are the following:

- The primary model data  $(A, B, C, \bar{\lambda}, w)$  is readily available in practice. This is because our partner’s warehouses use a sophisticated data collection system involving bar-code

scanners carried by pickers and packers and also placed in many locations in the conveyor system, induction stations and sorter chutes. This system generates a database of detailed flow timing information for individual orders from which our model’s service time distributions can be constructed;

- The primary model output  $(X(\tau), Y(\tau), Z(\tau))$  is directly observable in practice, and in fact that information corresponds exactly to that actually observed by managers when adjusting the revolving batch size. This model feature thus guarantees that the informational requirements of any closed loop release control policy developed will be realistic, and also allows for a quantitative model validation to be performed (see §3.2.2);
- Note that  $Y(\tau) + Z(\tau)$  represents at any time  $\tau$  the total number of busy chutes (either incomplete, complete and waiting for a packer, or being packed). The occurrence of gridlock can thus be directly expressed in terms of our model’s primary output as the event  $Y(\tau) + Z(\tau) > n$ , where  $n$  denotes the total number of sorter chutes. Besides, consider any release control policy  $\lambda$  ensuring with sufficiently high probability that this event does not occur (we discuss in §3.3 how this may be achieved). Under such policy the seemingly salient model assumption of infinite buffer size at the third station is in fact immaterial.

The appealing features of this model formulation do come with a price however. The actual time-to-chute and chute-dwell time of any particular order depend directly on the transit times between the picking area and the sorter of all the items it contains. In turn, those transit times are affected in practice by the congestion upstream of the sorter<sup>7</sup>. But the congestion upstream of the sorter is directly related to the output process  $X(\tau)$  representing the number of orders in the first station of our model. In summary, the service times  $A$  and  $B$  of the first two stations in our model could not a priori be considered exogenous, let alone stationary<sup>8</sup>. While one could conceivably attempt to create an analytical model predicting these service times as a function of the release rate and/or output process  $\{(X(\tau), Y(\tau), Z(\tau)) : \tau \geq 0\}$ , we believe that such model would be considerably more complicated than ours. To capture this endogeneity, we consider a small number of congestion levels  $g \in \{1, \dots, \bar{g}\}$  corresponding to adjacent consecutive ranges  $[d_g, d_{g+1})$  for conveyor system congestion, defined more precisely as the total number of items  $I(\tau)$  on the conveyor

---

<sup>7</sup> The Online Appendix contains a mathematical statement of the relationship between transit times and time-to-chute and chute-dwell time, as well as a study of the dependence of transit times on congestion.

<sup>8</sup> The same observation could conceivably be made for the service times of the third station because the walking time of packers also appears endogenous. However, data shows that  $C$  is in fact fairly stationary.

system between the picking area and the sorter. We then fit the distributions of  $A$  and  $B$  for all time periods in our data set when the system was in each congestion level, and thus obtain empirical distributions  $A(g)$  and  $B(g)$  for all the congestion levels  $g$  just defined. Remarkably, the distributions  $A(g)$  and  $B(g)$  constructed by this method are unimodal for each  $g$  (whereas the empirical distribution of  $A$  and  $B$  obtained for the entire data set are not), thus validating our approach at a qualitative level (see the Online Appendix for empirical density plots of these service time distributions).

Finally, the last modeling step consists of identifying a relationship between the total number of items  $I(\tau)$  on the conveyor system (which characterizes the congestion level  $g$ ) and the output  $(X(\tau), Y(\tau), Z(\tau))$ , so that the dynamics of our model be well defined. With  $\mathbb{E}[M]$  denoting the average number of items per customer order as before, the expression  $\mathbb{E}[M]X(\tau)$  results in a significant underestimation of the number of items  $I(\tau)$  in process between the picking area and the sorter, because many items still on the conveyors belong to customer orders with one or several other items already in a chute, and are therefore not accounted for by the process  $X$ . However, the expression  $I(\tau) \approx \mathbb{E}[M](X(\tau) + Y(\tau)/2)$  provides a relatively accurate estimate for the total number of items on the conveyors – this expression corresponds to the approximation that the order statistics of the arrival times of the items of each shipment to their corresponding chute are equally spaced in expectation. Ultimately, this method and the resulting predictive accuracy of our model are validated by performing a comparison of actual system output and simulated model output over time for given starting conditions and picking rate history (see §3.2.2).

## 3.2 Approximate Dynamics

**3.2.1 Derivation** Our next step is to study the dynamics of the queueing model described in §3.1, that is characterize how the process  $(X_t, Y_t, Z_t)$  evolves over time as a function of any release control policy  $\lambda$  considered. We want to understand in particular how such policy may dynamically affect the likelihood of the gridlock event  $\{Y_t + Z_t > n\}$ . Unfortunately, the exact analysis of that queueing model appears challenging because (i) the short control time period  $\delta$  precludes the use of any steady-state analysis; and (ii) service time distributions for the first two stations depend on the congestion level and are thus state-dependent. These observations motivate the development of an approximate version of our queueing model that is more amenable to analysis but still offer a suitably realistic representation of the actual

pick-to-ship process described earlier. We specifically consider the following approximations:

- *The second queue has an infinite number of servers.* This assumption is justified in light of the optimization problem that we consider eventually (in §3.3), where the probability of the gridlock event  $\{Y(\tau) + Z(\tau) > n\}$ , which contains the event  $\{Y(\tau) > n\}$ , is constrained to be very small.
- *The service times  $A(g)$ ,  $B(g)$  and  $C$  of the three queueing stations follow exponential distributions with first moments given by actual data.* The empirical distributions of  $B(g)$  and  $C$  constructed from data have coefficients of variation that are close to 1 and shapes that are similar to that of an exponential (see Online Appendix). However, the empirical distributions we constructed for  $A(g)$ , which have positive and relatively large support lower bounds determined by the speed of conveyors, do not. Note that Weber (2005) derives more accurate system dynamics for this model through an approximating queueing network that is still Markovian, using phase-type distributions and results on the  $M_t/G/\infty$  queue from Eick, Massey and Whitt (1993). However, we have found that the resulting increase of the number of state space dimensions does increase computational requirements for our approximate DP algorithm to a level that is impractical, at least at the time of writing.
- *Orders move at most one station downstream during each time period  $[t\delta; (t+1)\delta)$ .* This approximation substantially simplifies system dynamics, and does not seem to much harm prediction accuracy: because the actual expected service times  $\mathbb{E}[A(g)]$  and  $\mathbb{E}[B(g)]$  at the first and second stations are several times larger than the control period  $\delta$ , the transitions that this assumption ignores have very low probability relative to all others.
- *The congestion level remains constant within each control period  $[t\delta; (t+1)\delta)$ .* When simulating the exact queueing dynamics for the model described in §3.1 under various policies and input parameters, we have found that consecutive changes of congestion level occurring less than  $\delta$  time units apart were very rare.
- *The minimum of the numbers of packers  $w$  and closed chutes  $Z(\tau)$  remains constant within each control period  $[t\delta; (t+1)\delta)$ .* We have likewise observed that policies performing seemingly well in simulations resulted in a relatively high capacity utilization for the third queue, yielding  $\min(w, Z(\tau)) = w$  with high probability.

From elementary properties of markovian queues, the above approximations result in

the following discrete-time system dynamics (see Weber 2005 for this and other related derivations of transient system dynamics)<sup>9</sup>:

$$\left\{ \begin{array}{l} X_{t+1} = X_t + N_t^{\rightarrow X} - N_t^{X \rightarrow Y} \\ Y_{t+1} = Y_t + N_t^{X \rightarrow Y} - N_t^{Y \rightarrow Z} \\ Z_{t+1} = Z_t + N_t^{Y \rightarrow Z} - N_t^{Z \rightarrow} \\ g_t = \sum_{g=1}^{\bar{g}} g 1_{[d_g, d_{g+1})}(\mathbb{E}[M](X_t + \frac{Y_t}{2})) \end{array} \right. \quad \text{with} \quad \left\{ \begin{array}{l} N_t^{\rightarrow X} \sim \text{Poisson}(\lambda_t \delta) \\ N_t^{X \rightarrow Y} \sim \text{Binom}(X_t, 1 - e^{-\frac{\delta}{\mathbb{E}[A(g_t)]}}) \\ N_t^{Y \rightarrow Z} \sim \text{Binom}(Y_t, 1 - e^{-\frac{\delta}{\mathbb{E}[B(g_t)]}}) \\ N_t^{Z \rightarrow} \sim \text{Poisson}(\frac{(w \wedge Z_t) \delta}{\mathbb{E}[C]}) \end{array} \right. , \quad (1)$$

where the four random variables  $(N_t^{\rightarrow X}, N_t^{X \rightarrow Y}, N_t^{Y \rightarrow Z}, N_t^{Z \rightarrow})$  represent the number of customer orders that are respectively released into the first station, moved from the first to the second and the second to the third station, and processed out of the third station, between time periods  $t$  and  $t + 1$ . An appealing feature is that simulating system (1) only involves generating four standard random variables in each time period, and can thus be performed very efficiently. Computations can be even further reduced by substituting the binomial variables  $N_t^{X \rightarrow Y}$  and  $N_t^{Y \rightarrow Z}$  with normal random variables having the same mean and variance, which from the De Moivre-Laplace theorem is asymptotically exact for the large values of  $X_t$  and  $Y_t$  that are typical of our setting.

**3.2.2 Validation** Our next step was to validate the approximate queueing dynamics (1) using numerical simulation; the high-level procedure was to compare the predicted model state under some given release rate and packer staffing history against that actually observed in the real system when subjected to the same input. Specifically, we collected data series recording the actual state evolution  $(x^*(\tau), y^*(\tau), z^*(\tau))$ , actual control history  $\lambda^*(\tau)$  and actual number of staffed packers  $w^*(\tau)$ , each with one data point per minute and spanning a period of several days during a peak demand period faced by our industrial partner. In the context of the order release control problem that we formally define in the next section, a particularly relevant prediction lead-time is the control period  $\delta$  (of the order of a few minutes), since the associated dynamic program involves an expectation of the value function at time  $\tau + \delta$  given the system state at time  $\tau$ . We thus computed for every time  $\tau$  the average release rate over the following period of length  $\delta$ ,  $\tilde{\lambda}^*(\tau) = \frac{1}{\delta} \sum_{i=0}^{\delta-1} \lambda^*(\tau + i)$ , and

---

<sup>9</sup> In particular, the last expression of (1) corresponds to the departure process over a period of length  $\delta$  from a Markovian queue with  $w$  servers assumed to work continuously, except in the ramp-up periods. This assumption is appropriate given the very high utilization of that queue for most of the scenarios considered (particularly when nearing the gridlock state), and in the other scenarios we verified that the associated expression was still an acceptable approximation.



simulated the random variables  $(X_{t+1}, Y_{t+1}, Z_{t+1})$  characterized by (1) given  $(X_t, Y_t, Z_t) = ((x^*(\tau), y^*(\tau), z^*(\tau)), \lambda_t = \tilde{\lambda}^*(\tau)$  and  $w = w^*(\tau)$ . We then compared them with the actual state for the corresponding period  $(x^*(\tau+\delta), y^*(\tau+\delta), z^*(\tau+\delta))$ . Note that the historical data available only corresponds to a specific realization of our stochastic predictive model. For this reason, any discrepancy between the actual state  $(x^*(\tau+\delta), y^*(\tau+\delta), z^*(\tau+\delta))$  and (say) the estimated mean of the random variables  $(X_{t+1}, Y_{t+1}, Z_{t+1})$  just defined should be interpreted in light of the variability predicted by the model around those means. While our findings were consistent across all state variables, due to space constraints we only report here actual and predicted values for the number of busy chutes, which is particularly relevant in this context because the occurrence of gridlock is modeled by the event  $\{Y_t + Z_t > n\}$ . Specifically, Figure 1 shows the mean  $E[Y_{t+1} + Z_{t+1}]$  and associated centered empirical range with length  $6\sigma[Y_{t+1} + Z_{t+1}]$  thus estimated at each record time point  $(\tau)$  over one full (representative) day, along with the actual corresponding historical value  $y^*(\tau + \delta) + z^*(\tau + \delta)$ . Also highlighted in Figure 1 (with a gray background) are the time periods corresponding either to workers' breaks (from approximately 1:30 to 2:00, 5:30 to 6:00, 8:00, 10:15, 12:00 to 13:00, 15:15, 17:30 to 18:00, 20:15, 22:15, 23:30 onwards) or reduced activity due to shift change-over, equipment maintenance, breakdown or repair (around 0:15, 3:15, 11:30 to 12:00, 21:00).

Our main observation on the results shown in Figure 1 is that the time periods when the actual number of busy chutes falls outside of the empirical range predicted by our model coincide almost exactly with the workers' breaks and episodes of equipment maintenance/breakdown mentioned above. Furthermore, in all these periods the model significantly overestimates the number of busy chutes. This overestimation follows from the fact that our queueing model does not capture explicitly the induction stations at which items coming from the conveyor system are individually placed on the sorter's tilting trays. Indeed, the transition rate between the first and second stations in our queueing model only depends on the number of orders in the first station as well as its service time (time-to-chute) distribution, and thus does not directly account for the staffing of induction stations. This modeling choice is justified during the regular (non-highlighted) working hours, as the induction stations have appropriate processing capacity then. However, the periods highlighted in Figure 1 drastically impact the staffing of these induction stations, so that the actual flow of items into the sorter then is either considerably reduced (maintenance/breakdown) or stopped




Figure 1: Predicted Distribution  $Y_{t+1} + Z_{t+1}$  and Actual Value  $y^*(\tau + \delta) + z^*(\tau + \delta)$  of the Number of Busy Chutes over a 24 Hour Period.

(work breaks). Indeed, note that the actual number of busy chutes remains constant during all work breaks listed above, which reflects that actual flows into the sorter (induction) and out of it (packing) are stopped then. While the model correctly captures the packing rate reduction during such periods through its input data  $w$ , it ignores the corresponding decrease in induction rate, leading to the overestimation observed.

While the results observed during the highlighted periods enhanced our understanding of the relationship between our model and the actual system, they did not seem relevant to validation since gridlock may not occur during periods of forced reduced activity. During regular working hours the actual number of busy chutes observed almost always lied within our model's predicted range, so that the approximate dynamics tested appeared sufficiently accurate given our purposes; this validation exercise was thus deemed conclusive.

**3.3 Optimization Problem Formulation** We now state and discuss the formulation  $CDP[\beta]$  which provides the framework of our optimization study:

$$\begin{aligned}
CDP[\beta] : \quad & \max_{\lambda} \quad \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t \lambda_t | X_0, Y_0, Z_0] \\
\text{s.t.:} \quad & \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t 1_{\{Y_t^\lambda + Z_t^\lambda > n\}} | X_0, Y_0, Z_0] \leq \beta \\
& \lambda_t \in [0, \bar{\lambda}] \text{ for all } t \in \mathbb{N},
\end{aligned} \tag{2}$$

where  $\alpha \in (0, 1)$  is a discount factor,  $1_{\{\cdot\}}$  is an event indicator function,  $\beta$  is the *risk budget* or parameter defining the level of risk tolerated for the event of gridlock (see discussion below), expectations are taken over the sample space of release and service time realizations, and  $(X_0, Y_0, Z_0)$  is the initial state of the system. In (2), the maximum is taken over all stationary closed-loop and non-anticipative policies  $\lambda$ , and the notations  $Y_t^\lambda$  and  $Z_t^\lambda$  reflect the dependence of the model output process  $(Y_t, Z_t)$  on the release control policy considered. Since no ambiguity arises from the present context however, we will almost always omit that dependence in the following.

The objective function in (2) captures the goal of maximizing the throughput of the pick-to-ship process considered. Observe however that it is the (discounted) sum of release rates, which are proportional to the process input as opposed to process output – this is justified by the first constraint, which effectively prevents any unbounded accumulation of inventory in the system and is further discussed below. Note that the discount factor  $\alpha$  introduces a preference for units shipped in earlier periods. The classical interpretation of such discount factor as one minus a Bernoulli probability that the future stream of rewards may be interrupted is appealing here: when running into gridlock, the real pick-to-ship process goes through a lengthy recovery procedure which is not captured by the queueing model described in §3.1. In addition, we may consider  $\alpha$  for practical purposes as a tuning parameter affecting the features and performance of the policies derived. However, the primary reason for us to study here the discounted cost formulation (2) is that it is easier to solve numerically than the natural average cost formulation of the same problem. That latter formulation is formally linked to (2) through the following limiting statements<sup>10</sup>, which are proven in Blackwell (1962) and hold for any initial state:

$$\left\{ \begin{aligned}
& \lim_{\alpha \rightarrow 1^-} (1 - \alpha) \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t \lambda_t | X_0, Y_0, Z_0] = \lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E}[\sum_{t=0}^{k-1} \lambda_t] \\
& \lim_{\alpha \rightarrow 1^-} (1 - \alpha) \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t 1_{\{Y_t + Z_t > n\}} | X_0, Y_0, Z_0] = \lim_{k \rightarrow \infty} \frac{1}{k} \mathbb{E}[\sum_{t=0}^{k-1} 1_{\{Y_t^\lambda + Z_t^\lambda > n\}}] = \lim_{t \rightarrow \infty} \mathbb{P}(Y_t + Z_t > n)
\end{aligned} \right. \tag{3}$$

---

<sup>10</sup> Blackwell proves the first equality in each line for a finite state Markov Chain. The second equality in the second line follows from the ergodicity of Markov chains with unique steady-state distributions.

The second and third equality statements in (3) imply that the first constraint in (2) is asymptotically equivalent as  $\alpha \rightarrow 1$  (a very relevant regime given the numerical value we use for  $\alpha$  are very close to 1) to the much more intuitive expression

$$\lim_{t \rightarrow \infty} \mathbb{P}(Y_t + Z_t > n) \leq (1 - \alpha)\beta,$$

which specifies an upper bound on the steady-state probability that the system is in a state of gridlock; the comparably unintuitive exact expression of that constraint in (2) (i.e. a discounted sum of indicator functions) is only motivated by technical dynamic programming considerations (see §3.4). From a modeling perspective, that constraint thus balances the throughput maximization objective in (2) with the need to avoid the state of gridlock. Note that it is only a probabilistic statement: because the support of the inter-arrival and service time distributions we use are neither bounded from above or bounded away from zero, with any policy resulting in some positive release rates it is impossible to guarantee in a deterministic sense that gridlock will never occur.

Finally, observe that the statement of problem  $CDP[\beta]$  depends on the initial state  $(X_0, Y_0, Z_0)$ . However, we have used values of the discount factor  $\alpha$  that are very close to 1 in our experiments, and observed that the choice of the initial state had very little impact on the results, if any. This is explained in part by the limiting statements (3), where the r.h.s is independent of the initial state, as is typical of the objective of an average cost DP formulation.

**3.4 Optimization Algorithm** We have specifically formulated the dynamic program  $CDP[\beta]$  in (2) so it would belong to a family of constrained Markov decision processes for which some theoretical results and approximate computational methods can be easily derived (see Altman 1999 for a review, and Hordijk and Spieksma 1989 for an application of these methods to a queueing model with features similar to ours). In particular, we now outline a method for computing a solution to  $CDP[\beta]$  by solving a sequence of related unconstrained dynamic programs  $UDP[\theta]$  obtained for any  $\theta \geq 0$  as

$$\begin{aligned} UDP[\theta] : \quad & \max_{\lambda} \quad \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t (\lambda_t - \theta \cdot 1_{\{Y_t + Z_t > n\}}) | (X_0, Y_0, Z_0) = (x, y, z)] \\ & \text{s.t.:} \quad \lambda_t \in [0, \bar{\lambda}] \text{ for all } t \in \mathbb{N}, \end{aligned} \tag{4}$$

where the underlying state dynamics are identical to those of the original problem  $CDP[\beta]$ . The problem  $UDP[\theta]$  just defined is thus a Lagrangian relaxation of  $CDP[\beta]$  where the

first constraint in (2) is now captured through the objective function and weighted there by the multiplier  $\theta$ , to be interpreted as an instant penalty for entering a gridlock state, i.e.  $\{Y_t + Z_t > n\}$ . Define now  $\mathbf{j}^\theta(x, y, z)$  as the optimal cost-to-go function for  $UDP[\theta]$ , equal to  $\mathbf{r}^\theta(x, y, z) - \theta \cdot \mathbf{c}^\theta(x, y, z)$  with  $\mathbf{r}^\theta(x, y, z) \triangleq \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t \lambda_t^\theta | (X_0, Y_0, Z_0) = (x, y, z)]$  and  $\mathbf{c}^\theta(x, y, z) \triangleq \mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t 1_{\{Y_t^{\lambda^\theta} + Z_t^{\lambda^\theta} > n\}} | (X_0, Y_0, Z_0) = (x, y, z)]$ , where  $\lambda^\theta$  is an optimal policy for  $UDP[\theta]$ . The following results are obtained through a straightforward adaptation to the discounted case of the proofs of Lemma 3.1, Theorem 4.3 and Theorem 4.4 from Beutler and Ross (1985) and Corollary 3.5 from Beutler and Ross (1986):

**Lemma 1** *There exists a stationary optimal policy  $\lambda$  for  $CDP[\beta]$  that is deterministic in all states but one, and randomizes between at most two actions in that state. Moreover,  $\lambda$  achieves  $\mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t 1_{\{Y_t + Z_t > n\}} | (X_0, Y_0, Z_0) = (x, y, z)] = \beta$  and there exists  $\theta^* \geq 0$  such that  $\lambda$  is optimal for  $UDP[\theta^*]$ .*

**Lemma 2** *Suppose that for some  $\theta \geq 0$  there exists a policy  $\lambda^\theta$  such that  $\lambda^\theta$  is optimal for  $UDP[\theta]$  and achieves  $\mathbb{E}[\sum_{t=0}^{+\infty} \alpha^t 1_{\{Y_t + Z_t > n\}} | (X_0, Y_0, Z_0) = (x, y, z)] = \beta$ . Then  $\lambda^\theta$  is optimal for  $CDP[\beta]$ .*

**Lemma 3** *For any initial state  $(x, y, z)$ ,  $\mathbf{j}^\theta(x, y, z)$ ,  $\mathbf{r}^\theta(x, y, z)$  and  $\mathbf{c}^\theta(x, y, z)$  are all monotone non-increasing in  $\theta$ .*

The solution method we have implemented consists of a line search over  $\theta$ , where the optimal solution  $\lambda^\theta$  to  $UDP[\theta]$  is computed at each iteration along with the corresponding cost-to-go functions  $\mathbf{j}^\theta$ ,  $\mathbf{r}^\theta$  and  $\mathbf{c}^\theta$  using standard approximate DP methods (see below), and the search proceeds until a value of  $\theta$  achieving  $\mathbf{c}^\theta(x, y, z) \approx \beta$  is found. Lemma 1 asserts that such  $\theta$  exists; Lemma 2 suggests that once such  $\theta$  is found, the resulting policy  $\lambda^\theta$  should be (near) optimal for  $CDP[\beta]$ ; finally the monotonicity of  $\mathbf{c}^\theta$  shown in Lemma 3 indicates that an efficient search can be used. The specific algorithm we have implemented is a dichotomic search over a specified interval  $[\underline{\theta}, \bar{\theta}]$ , with an accuracy termination parameter  $\epsilon$ . While a more detailed description and discussion of convergence properties can be found in the Online Appendix, we observe here that this algorithm may require to solve up to  $\log_2(\frac{\bar{\theta} - \underline{\theta}}{\epsilon})$  unconstrained dynamic programs  $UDP[\theta]$  in order to compute a solution to  $CDP[\beta]$ . The associated computational efforts may thus seem daunting at first. However, we were able to reduce computations to a practical level through the following additional steps: (i) In order to solve each instance of  $UDP[\theta]$ , we use an approximate policy iteration algorithm relying

on the Robbins-Monro stochastic approximation scheme for the evaluation step, and Monte-Carlo simulations for the improvement step (see the Appendix for a detailed statement of this algorithm, related references and discussion); and (ii) At the  $k$ -th iteration of the search algorithm (with corresponding multiplier value  $\theta^k$ ), we use the best policy found for  $UDP[\theta^{k-1}]$  at the previous iteration as a starting point to the policy iteration algorithm used to solve problem  $UDP[\theta^k]$ .

In the remainder of this paper, we refer to the policy obtained from the algorithm just stated as  $ADP^\beta$  (the superscript  $\beta$  is omitted when no ambiguity arises) and denote its release rate function as  $\lambda^{ADP}(x, y, z)$  or  $\lambda_t^{ADP} \triangleq \lambda^{ADP}(X_t, Y_t, Z_t)$ .

**3.5 Policy Structure** We now discuss the qualitative features of policy  $ADP^\beta$ . Theoretical results on the structure of optimal policies for problem (2) have so far eluded us, which is relatively unsurprising given the relative complexity of the underlying model (a queueing network with discrete-time controls and service times depending on a function of congestion in some parts of the network)<sup>11</sup>. As a result, the following discussion is based instead on a large number of consistent empirical observations.




Figure 2: Release Rate Function of Policy  $ADP^\beta$  for the Scenarios  $w \in \{0.75p, p, 1.25p\}$  and  $\beta(1 - \alpha) = 10^{-6}$ .

The main observed features of policy  $ADP^\beta$  are illustrated by Figure 2, which includes plots of its normalized release rate  $\lambda_t^{ADP}$  as a function of the reduced and normalized state

<sup>11</sup> While it is conceivable that such results could be obtained assuming simplified model dynamics (with exogenous service times for example), such a study falls outside the scope of the present paper, in part because validation experiments of the type described in §3.2.2 show that simplified versions of the model we consider have poor predictive accuracy in the industrial setting motivating our study.

$(X_t/n, (Y_t + Z_t)/n)^{12}$  for three representative scenarios characterized by a limiting gridlock probability of  $10^{-6}$  and a number of packers  $w \in \{0.75p, p, 1.25p\}$ , where  $p$  is the average number of packers used by our industrial partner during peak demand periods. A first intuitive feature observed is that the release rate is a decreasing function of the state, in the sense that  $\lambda^{ADP}(x', y', z') \leq \lambda^{ADP}(x, y, z)$  when  $x' \geq x$ ,  $y' \geq y$  and  $z' \geq z$ . In particular, the policy releases orders at the maximum rate allowed when the system is almost empty (this region of the state space corresponds to the upper middle corner of all plots in Figure 2), and stops releases altogether in the regions of the state-space corresponding to heavy congestion (as seen in the lower middle corner of the plots). A second noteworthy feature is that the release rate function displays sudden drops, as seen around  $X_t/n \approx 85\%$  in Figure 2 (a),  $X_t/n \approx 95\%$  in Figure 2 (b) and  $X_t/n \approx 100\%$  in Figure 2 (c). While this is not obvious from these figures alone because of the state reduction used, it can be verified that these drops correspond to transitions of the congestion level  $\mathbb{E}[M](X_t + Y_t/2)$  between two consecutive ranges of values characterizing system dynamics (see §3.2.1). Indeed, simulation shows that these drops enable the policy to maintain the system in a desirable steady-state congestion level<sup>13</sup>.

Finally, the structure of policy  $ADP^\beta$  is sensitive to the number of packers. Figure 2 (a) shows that with a small number of packers, the release rate function  $\lambda_t^{ADP}$  is almost independent of  $X_t$ . In such a regime good policies must heavily utilize packing capacity, and any given change in the arrival rate of orders to the third packing station has a substantial impact on its occupancy process, or number of green chutes  $Z_t$ . Consequently, the  $ADP^\beta$  policy thus reacts considerably more to changes in the state variable  $Z_t$  than to changes of  $X_t$  or  $Y_t$ , which is easily verified quantitatively (along with all similar statements in this discussion) by examining the relative values of the partial derivatives  $\frac{\partial \lambda^{ADP}(x,y,z)}{\partial x}$ ,  $\frac{\partial \lambda^{ADP}(x,y,z)}{\partial y}$  and  $\frac{\partial \lambda^{ADP}(x,y,z)}{\partial z}$  in relevant regions of the state space. In addition,  $ADP$  compensates then for even small deviations around an implicit target value for  $Z_t$  with drastic changes in its instantaneous release rate, as illustrated by the sudden drop of the release rate surface seen in Figure 2 (a) around  $(Y_t + Z_t)/n \approx 80\%$ . In contrast, Figure 2 (c) illustrates that with a high number of packers, policy  $ADP^\beta$  reacts much more to changes in the number

---

<sup>12</sup> More precisely the function plotted in the plane  $(X_t, Y_t + Z_t)$  is  $f(x, b) \triangleq \frac{1}{b+1} \sum_{j=0}^b \lambda^{ADP}(x, j, b-j)$ .

<sup>13</sup> The Online Appendix contains a more extensive discussion on why some congestion levels are preferable to others.

of orders in transit  $X_t$  than to changes in the number of busy chutes  $Y_t + Z_t$ , and it can be verified that  $\lambda_t^{ADP}$  is in fact almost independent of  $Z_t$  then. In this scenario with low packing utilization, completed chutes tend to be attended to immediately by a packer, i.e. the jobs representing them in our model experience little or no queueing in the third station. Consequently, any temporary increase of  $Z_t$  around its operating steady-state average is absorbed by spare packing capacity, and likely corrected by the time any change in the order release rate can have any impact on the sorter (second and third queueing stations), as the expected time-to-chute  $\mathbb{E}[A(g)]$  is long relative to the pack-to-pack time  $\mathbb{E}[C]$  (see §3.1).

## 4 Wave Release Models

The model described in the previous section was primarily constructed with the goal of helping our industrial partner optimize its waveless release policy. The present section discusses how it can also be modified in order to support a partial performance comparison with wave-based policies through simulation. As observed by Johnson and Lofgren (1994) and others, these more traditional policies effectively decouple the warehouse areas of picking and sorting, since only batches of orders that have been completely picked (the waves) are typically released into the sorter. Our wave release models reflect this decoupling, in that they only consider the picking operation and the conveyor system leading to the sorting area through the assumption that complete waves of picked orders are always ready to be released into the sorter for induction, sorting and packing. This approach enables a meaningful comparison between wave-based and waveless release policies along the performance dimensions of throughput (the primary concern of our partner during the peak demand periods which motivated our study), gridlock probability, packer utilization and sorter utilization. However, it leaves aside the dimensions of order cycle time as well as storage and/or conveyor space requirements for the intermediary buffer between picking and sorting, which as emphasized in §1 constitute substantial negatives of wave picking. Finally, picker utilization is another important performance dimension that we are unable to fully explore, because our models do not explicitly capture the detailed layout and resulting picking tours in the picking area (see §3.1 and related discussion in §6).

In order to enable a meaningful comparison with waveless policies, we consider models capturing the most sophisticated wave release policies observed in practice (see §1 for background), as described in §4.1 and §4.2 below.



**4.1 Overlapping Waves Policy** Under this policy and according to common industrial practice (see §1), complete waves of items corresponding to a number of orders equal to the number of sorter chutes ( $n$ ) are successively released into the sorter. This is modeled using a three station serial queueing network sharing several features with the one stated in §3.1 and defined as follows. First, each wave is modeled as a sequence of items generated by simulating  $n$  independent draw from a distribution  $M$  of the number of items per order constructed from empirical data, and assuming that the items of each order are uniformly distributed within the wave (as in Hinojosa 1996 and Johnson 1998). Also, the first and last items of each order within the wave are tagged<sup>14</sup>, for a reason that will soon become clear. Note that in the waveless picking model of §3 induction capacity is only considered implicitly through the dependence of transit times on the congestion upstream of the sorter. While justified for the stationary steady-state associated with waveless release, that approach is not appropriate to model transient wave-based release. This is because induction stations are periodically faced then with the sudden release of large batches of orders, so that the congestion generated by their capacity limitations does becomes material over short time periods. The first station in our wave picking model thus represents the induction stations, and it processes individual items from each incoming wave with a service time retrieved from the actual dataset of flow information obtained from our partner. The second station is an infinite server queue representing as before the incomplete chutes, however its service time for each order is now given endogenously by the time between the completion of its first and last items by the induction station (hence the tagging mentioned above). Finally, the third station represents the packing queue and is identical to that described in §3.1.

In the simplest form of wave picking, each wave is released into the sorter just when the last order of the previous one is packed. As in Bozer, Quiroz and Sharp (1988) and Johnson and Lofgren (1994) however, we consider the more general release policy whereby each wave is released as a given percentage of the chutes in the sorter (denoted  $\Omega \in (0, 100]$ ) become empty. This policy is easily simulated using the model defined above, and is referred to as  $W_\Omega$  in the remainder of this paper. Note that the simple non-overlapping policy described earlier corresponds to the particular case  $W_{100}$ . Also, overlapping waves ( $\Omega < 100$ ) give rise to the possibility of gridlock, which is still characterized by the event  $Y(\tau) + Z(\tau) > n$ , with

---

<sup>14</sup> The waves do not include single item orders, which are packed through a separate process in the warehouses of our industrial partner.

$Y(\tau)$  and  $Z(\tau)$  representing as before the number of orders in the second and third stations at time  $\tau$ , respectively. This leads us to define  $\Omega^\beta$  as the wave overlapping parameter such that the corresponding policy  $W_{\Omega^\beta}$  achieves the same steady-state gridlock probability as policy  $ADP^\beta$ , i.e.  $\lim_{\tau \rightarrow \infty} \mathbb{P}(Y(\tau) + Z(\tau) > n) = (1 - \alpha)\beta$  (see §3.3).

**4.2 Split Sorter Policy** As described in §1 and Ruben and Jacobs (1992), Russell and Meller (2003) and Perkins (2008), this alternative policy involves the sequential release of waves with a number of orders equal to half the number of chutes in the sorter, or  $n/2$ . The sorter chutes are split in two halves (each dedicated to a separate wave), packers give priority to completed orders in the oldest wave, and a new wave is released as soon as one of the two halves becomes empty. The model we use to represent this policy is the same as described in §4.1, except that the size of each wave is halved and the second station as well as the queue of the third station are duplicated, creating a fork from the first station and a merge into the servers of the third station. Each wave is assigned the second station duplicate which was empty upon its release, and this assignment is implemented at the fork following the first (induction) station. Finally, orders belonging to the oldest wave are given a higher priority by packers. That is, packers may start working on orders from a more recent wave, but only if no order from the previous wave is ready to be packed. In the remainder of this paper, this policy is referred to as  $W/2$ .

## 5 Numerical Experiments

The goal of our simulation study is to estimate and understand the relative performance of various order release control policies in the setting of our industrial partner’s warehouses. We first review the policies considered and other experimental design issues in §5.1, then present and discuss our results in §5.2. The Online Appendix also contains additional experiments designed to assess the robustness of the waveless release policies considered with respect to transient disruptions.

**5.1 Experimental Design** In addition to the waveless release policy  $ADP^\beta$  derived in §3 and the wave-based release policies  $W_\Omega$  and  $W/2$  defined in §4.1 and §4.2 respectively, we also consider the following simple waveless release policies:

**Policy  $CST^\beta$  (constant release):** Releases orders at the constant rate  $\lambda^{CST} \in [0, \bar{\lambda}]$  corresponding to the best constant solution to (2). That rate is easily found by simulation-based

search.

**Policy  $CWP^\beta$  (constant work in process):** Releases orders at a rate given by the function

$$\lambda^{CWP}(x, y, z) = \begin{cases} \bar{\lambda}^{CWP} & \text{if } x + y + z < \bar{k} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where the parameters  $\bar{\lambda}^{CWP} \in [0, \bar{\lambda}]$  and  $\bar{k} \in \mathbb{N}$  are likewise determined by simulation-based search so that the resulting policy is the best solution to (2) within the family of CONWIP policies defined by (5). We consider CONWIP policies here because despite their simplicity, they have been found to perform well in many environments (Spearman and Zazanis 1992).

The range of simulation scenarii we consider is characterized by three different values for the number of packers  $w \in \{p, 1.125p, 1.25p\}$ , where  $p$  denotes the average number of packers working in our industrial partner's warehouse during a peak demand period, under the staffing policies prevailing at the beginning of our interaction. We also consider two risk values  $\bar{\beta}$  (high risk) and  $\underline{\beta}$  (low risk), which correspond under our assumed discount factor  $\alpha = 0.97$  to limiting gridlock probabilities of  $\bar{\beta}(1 - \alpha) \simeq 10^{-3}$  and  $\underline{\beta}(1 - \alpha) \simeq 10^{-6}$  (see §3.3). In practice the level of gridlock risk associated with  $\bar{\beta}$  is deemed unacceptably high by our industrial partner, but we consider it here to perform an analysis of sensitivity relative to the risk parameter. The main performance measure we investigate is the simulated *throughput*  $\gamma^D \triangleq \mathbb{E}[\sum_{t=0}^{k-1} \lambda_t^D]/k$  of each policy  $D \in \{ADP^\beta, CWP^\beta, CST^\beta, W_\Omega, W/2\}$ , where the notation  $\lambda_t^D$  denotes the simulated release rate of policy  $D$  at time  $t$  and time index  $k$  corresponds to 3.5 simulated days (graphical representations of system behavior suggest that all policies have long reached steady-state by then)<sup>15</sup>. In order to enable a meaningful assessment and preserve our partner's confidential information, all throughput results are provided as a ratio to the average throughput  $\gamma^{HIST}$  observed in our industrial partner's warehouse during a period with no breaks and  $p$  packers assigned to the sorter. Note that for all policies the packing capacity  $w(\mathbb{E}[C])^{-1}$  (or service rate of the third station described in §3.1 and §4) constitutes an upper bound for the throughput rate, and that the *packing utilization* is given by  $\gamma^D / (w(\mathbb{E}[C])^{-1})$ . While policies  $ADP^\beta$ ,  $CWP^\beta$ ,  $CST^\beta$  and  $W_{\Omega^\beta}$  are constructed so that their limiting gridlock probability is set by design to  $\beta(1 - \alpha)$ , we also report the estimated gridlock probability  $\mathbb{P}(\text{Gridlock})$  associated with  $W_\Omega$  for other values of  $\Omega$ . Finally, we also report the *sorter utilization* given by  $\mathbb{E}[Y_\infty^D + Z_\infty^D]/n$ , where the numerator is

---

<sup>15</sup> Note that  $\gamma^{CST} = \lambda^{CST}$ .

	Number of Packers ( $w$ )	$p$	$1.125p$	$1.25p$
	Chutes per Packer ( $n/w$ )	55	48.8	44
$ADP^\beta$ [ $ADP^{\bar{\beta}}$ ]	Throughput	104.4 [104.4]	109.3 [109.5]	110.3 [110.4]
	Packing Utilization	99.4 [99.4]	92.6 [92.8]	84.1 [84.1]
	Sorter Utilization	85.5 [88.2]	78.2 [79]	80 [81.2]
$CWP^\beta$ [ $CWP^{\bar{\beta}}$ ]	Throughput	104 [104.3]	105.3 [106.4]	105.8 [106.5]
	Packing Utilization	99.0 [99.3]	89.2 [90.1]	80.6 [81.2]
	Sorter Utilization	80.4 [83.2]	77.1 [78.2]	77.1 [78.9]
$CST^\beta$ [ $CST^{\bar{\beta}}$ ]	Throughput	101.9 [103]	102.7 [104.6]	102.8 [104.7]
	Packing Utilization	97.0 [98.0]	87.0 [88.7]	78.3 [79.8]
	Sorter Utilization	74.3 [74.5]	74.1 [74.3]	74.9 [75.2]
$W_{\Omega^\beta}$ [ $W_{\Omega^{\bar{\beta}}}$ ]	Throughput	98.7 [101.5]	108.9 [109.9]	110.3 [110.6]
	Packing Utilization	94.7 [97.2]	92.6 [93.5]	84.0 [84.6]
	Sorter Utilization	78.1 [80.7]	79.9 [80.9]	76.7 [76.8]
	$\Omega^\beta$ [ $\Omega^{\bar{\beta}}$ ]	49 [46]	44 [37]	31 [30]
$W_{60}$	Throughput	89.9	95.3	100
	Packing Utilization	86	81	76.5
	Sorter Utilization	70.7	69.9	69.3
	$\mathbb{P}(\text{Gridlock})$	0	0	0
$W_{100}$	Throughput	66.2	71.3	75.9
	Packing Utilization	63.3	60.6	52.5
	Sorter Utilization	52.1	52.3	52.6
	$\mathbb{P}(\text{Gridlock})$	0	0	0
$W/2$	Throughput	103.2	105.3	105.4
	Packing Utilization	98.7	89.6	80.7
	Sorter Utilization	61.6	38.2	36.2
	$\mathbb{P}(\text{Gridlock})$	0	0	0

Table 1: Numerical Simulation Results. Notes: All numbers shown in the third and subsequent rows are percentages. The length of the 95% confidence interval for all simulation results reported is smaller than 0.2% of the corresponding estimate.

an estimate obtained from simulation for the average number of busy chutes in steady-state under each policy  $D$  considered.

## 5.2 Results and Discussion

Table 1 summarizes our steady-state simulation results.

We discuss here separately the results for the waveless policies (in §5.2.1) then the results for the wave-based policies (in §5.2.2). A relative performance comparison between the two types of policies based on these results is provided as part of our concluding remarks in §6.

We also refer the reader to the Online Appendix for a discussion of additional simulation experiments conducted to assess the robustness of some of the policies considered with respect to transient shocks and inaccurate input data.

### 5.2.1 Waveless Release Policies

Table 1 shows that for the waveless policies  $D \in \{ADP^\beta, CWP^\beta, CST^\beta\}$  considered, the effective packing capacity utilization  $\gamma^D/w(\mathbb{E}[C])^{-1}$

is quite high at 97% and above when the number of packers  $w$  is equal to  $p$  (and, a fortiori, when  $w < p$ ), then drops to 92.6% and below for  $w = 1.125p$  and even more drastically at 84.1% and below for  $w = 1.25p$ . We observe that three factors may conceptually constrain throughput in this system: the maximum release rate  $\bar{\lambda}$ , the packing capacity  $w(\mathbb{E}[C])^{-1}$  and the gridlock probability constraint. Because the maximum release rate is substantially higher than the packing capacity, both in practice and in all simulation scenarios considered here, only the last two are relevant. With a relatively low number of packers  $w \leq p$ , packing is effectively the system bottleneck as the throughput of all policies remains relatively close then to the overall packing capacity. Because packing capacity is an upper bound on the long-term average throughput of all policies independently of the gridlock risk  $\beta$ , this also indicates that all three policies are near-optimal then, and that the gridlock probability constraint results in very little throughput loss relative to the unconstrained problem. When the number of packers increases ( $w > p$ ) however, both their effective utilization and the marginal gain in throughput from this additional packing capacity decrease under all policies considered, so that the gridlock constraint becomes the system bottleneck.

A deeper interpretation of these results stems from Theorem 1 in Chao and Scott (2000), which states that the stochastic processes representing the number of jobs in a set of  $G/M/w$  queueing systems with constant service effort  $w(\mathbb{E}[C])^{-1}$  increase with the number of servers  $w$  for the stochastic ordering relationship. This implies in our setting that the fractiles of the distribution of busy chutes  $Y_t + Z_t$  increase with the number of packers  $w$  when the overall packing utilization is held constant, or equivalently that with more packers a lower utilization is required to maintain any of these fractiles at a constant value (as the gridlock probability constraint requires). Another relevant insight from queueing theory is that the performance measures of highly congested queues are much more sensitive to a given change in their capacity utilization than that of less congested queues. Consequently, when the number of packers is low and packing utilization is high, even a small change in the release rate significantly impacts the fractiles of the distribution of busy chutes and the probability of gridlock. Equivalently, a given increase in the tolerated probability of gridlock affords little additional throughput then. Indeed, for every policy considered in Table 1 the additional average throughput obtained by increasing the gridlock risk parameter from  $\underline{\beta}$  to  $\bar{\beta}$  increases with the number of packers  $w$ , and it is almost negligible for policies  $ADP^\beta$  and  $CWP^\beta$  in the

high congestion scenario where  $w = p$ . This also explains why the throughput superiority of  $ADP^\beta$  relative to  $CWP^\beta$  increases from 0.4% for  $w = p$  to 4.5% for  $w = 1.25p$ , and that of  $CWP^\beta$  relative to  $CST^\beta$  increases from 2.1% to 3% as  $w$  increases from  $p$  to  $1.25p$  (similar results are observed with  $\beta = \underline{\beta}$ ). Indeed, the range of instantaneous release rates that do not lead to a violation of the gridlock probability constraint is more limited when  $w \leq p$  and packing utilization is high. As a result, the greater structural ability of  $ADP$  relative to  $CWP$  (resp.  $CWP$  relative to  $CST$ ) to dynamically adapt the release rate to process conditions does not provide substantial benefits then. As seen in Figure 3 however, when  $w = 1.25p$  policy  $ADP$  and to a slightly lesser extent  $CWP$  are more able to address temporary stochastic increases of the number of busy chutes  $Y_t + Z_t$  above their operating averages by reducing the instantaneous release rate accordingly, which results in a smaller volatility of the process  $Y_t + Z_t$ , and ultimately maintains higher sorter and packing utilization than  $CST$ , and therefore greater throughput, for the same level of risk.




Figure 3: Two Standard Deviation Range of the Steady-State Number of Busy Chutes  $\mathbb{E}[Y_t + Z_t] \pm 2\sigma[Y_t + Z_t]$  for Policies  $ADP^\beta$ ,  $CWP^\beta$  and  $CST^\beta$  with  $w = 1.25p$  Packers. Note: Statistics are computed after 2 days of simulated time.

We also believe that the significant decrease of packing capacity utilization just beyond the average number of packers  $p$  actually used by our industrial partner is not coincidental, and in fact lends support to the validity of our results. The organizational structure used in our partner's warehouse defines a picking team and a packing team, with separate perfor-

mance metrics. A key metric used for the packing team is the average worker productivity, defined for a given period of time as the number of orders packed divided by the corresponding number of man×hours used. This metric thus creates a local incentive to maximize packing capacity utilization, which explains that the actual staffing level of packers coincides with the point beyond which the marginal (throughput) return of additional packing capacity starts to markedly decrease. However, we submit that the appropriate staffing level of packers should be determined by weighing the labor cost incurred against the overall system throughput it enables. In particular, during peak demand periods when the financial benefits of additional throughput and short customer lead-times are particularly high, keeping packing capacity heavily utilized may not be as important per se, and the current policy may result in staffing less packers than is optimal, as surge packing capacity play an important role in avoiding gridlock. In that respect, the results shown in Table 1 should enable a more precise examination of this trade-off by our industrial partner, and a better understanding of the impact of local staffing policies on system throughput.

Finally, the results shown in Table 1 suggest that, in the case where  $p$  packers are assigned to the sorter, policies  $CST^\beta$ ,  $CWP^\beta$  and  $ADP^\beta$  may yield a throughput increase of 1.9% to 4.4% relative to the throughput  $\gamma$  observed in our industrial partner’s warehouse under comparable staffing conditions. It may be surprising at first that even policy  $CST^\beta$  outperforms the policy used in our partner’s warehouse. We point out however that despite its simplicity  $CST^\beta$  is still obtained through optimization (over the constant release rate  $\lambda^{CST}$ ), while the release policy used by our partner at the beginning of our interaction was not fully formalized and relied at least in part on the judgment of employees having sometimes little or no experience with warehouse dynamics during peak demand periods. Unfortunately, similar historical performance data was not readily available to us for numbers of staffed packers that are different than  $p$ . However, assuming that the relative performance of  $CST^\beta$  and our partner’s historical policy would be maintained in such scenarii, we can speculate from Table 1 that policy  $ADP^\beta$  (resp.  $CWP^\beta$ ) would only yield a very modest throughput improvement with fewer packers than  $p$ , but an increase in throughput close to 8% (resp. 3%) with 25% more packers than when  $w = p$ . In any case, our model predicts that the combined use of policy  $ADP^\beta$  and addition of 25% more packers than  $p$  would increase throughput relative to the historical performance we have observed by about 10%.























