



EXCERPT FROM THE PROCEEDINGS

OF THE
TENTH ANNUAL ACQUISITION
RESEARCH SYMPOSIUM
LOGISTICS MANAGEMENT

Wave Release Strategies to Improve Service in Order Fulfillment Systems

**Erdem Çeven and Kevin Gue
Auburn University**

Published April 1, 2013

Approved for public release; distribution is unlimited.
Prepared for the Naval Postgraduate School, Monterey, CA 93943.

Disclaimer: The views represented in this report are those of the authors and do not reflect the official policy position of the Navy, the Department of Defense, or the federal government.



The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Business & Public Policy at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).



ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Preface & Acknowledgements

Welcome to our Tenth Annual Acquisition Research Symposium! We regret that this year it will be a “paper only” event. The double whammy of sequestration and a continuing resolution, with the attendant restrictions on travel and conferences, created too much uncertainty to properly stage the event. We will miss the dialogue with our acquisition colleagues and the opportunity for all our researchers to present their work. However, we intend to simulate the symposium as best we can, and these *Proceedings* present an opportunity for the papers to be published just as if they had been delivered. In any case, we will have a rich store of papers to draw from for next year’s event scheduled for May 14–15, 2014!

Despite these temporary setbacks, our Acquisition Research Program (ARP) here at the Naval Postgraduate School (NPS) continues at a normal pace. Since the ARP’s founding in 2003, over 1,200 original research reports have been added to the acquisition body of knowledge. We continue to add to that library, located online at www.acquisitionresearch.net, at a rate of roughly 140 reports per year. This activity has engaged researchers at over 70 universities and other institutions, greatly enhancing the diversity of thought brought to bear on the business activities of the DoD.

We generate this level of activity in three ways. First, we solicit research topics from academia and other institutions through an annual Broad Agency Announcement, sponsored by the USD(AT&L). Second, we issue an annual internal call for proposals to seek NPS faculty research supporting the interests of our program sponsors. Finally, we serve as a “broker” to market specific research topics identified by our sponsors to NPS graduate students. This three-pronged approach provides for a rich and broad diversity of scholarly rigor mixed with a good blend of practitioner experience in the field of acquisition. We are grateful to those of you who have contributed to our research program in the past and encourage your future participation.

Unfortunately, what will be missing this year is the active participation and networking that has been the hallmark of previous symposia. By purposely limiting attendance to 350 people, we encourage just that. This forum remains unique in its effort to bring scholars and practitioners together around acquisition research that is both relevant in application and rigorous in method. It provides the opportunity to interact with many top DoD acquisition officials and acquisition researchers. We encourage dialogue both in the formal panel sessions and in the many opportunities we make available at meals, breaks, and the day-ending socials. Many of our researchers use these occasions to establish new teaming arrangements for future research work. Despite the fact that we will not be gathered together to reap the above-listed benefits, the ARP will endeavor to stimulate this dialogue through various means throughout the year as we interact with our researchers and DoD officials.

Affordability remains a major focus in the DoD acquisition world and will no doubt get even more attention as the sequestration outcomes unfold. It is a central tenet of the DoD’s Better Buying Power initiatives, which continue to evolve as the DoD finds which of them work and which do not. This suggests that research with a focus on affordability will be of great interest to the DoD leadership in the year to come. Whether you’re a practitioner or scholar, we invite you to participate in that research.

We gratefully acknowledge the ongoing support and leadership of our sponsors, whose foresight and vision have assured the continuing success of the ARP:



- Office of the Under Secretary of Defense (Acquisition, Technology, & Logistics)
- Director, Acquisition Career Management, ASN (RD&A)
- Program Executive Officer, SHIPS
- Commander, Naval Sea Systems Command
- Program Executive Officer, Integrated Warfare Systems
- Army Contracting Command, U.S. Army Materiel Command
- Office of the Assistant Secretary of the Air Force (Acquisition)
- Office of the Assistant Secretary of the Army (Acquisition, Logistics, & Technology)
- Deputy Director, Acquisition Career Management, U.S. Army
- Office of Procurement and Assistance Management Headquarters, Department of Energy
- Director, Defense Security Cooperation Agency
- Deputy Assistant Secretary of the Navy, Research, Development, Test, & Evaluation
- Program Executive Officer, Tactical Aircraft
- Director, Office of Small Business Programs, Department of the Navy
- Director, Office of Acquisition Resources and Analysis (ARA)
- Deputy Assistant Secretary of the Navy, Acquisition & Procurement
- Director of Open Architecture, DASN (RDT&E)
- Program Executive Officer, Littoral Combat Ships

James B. Greene Jr.
Rear Admiral, U.S. Navy (Ret.)

Keith F. Snider, PhD
Associate Professor



Logistics Management

Fully-Burdened Cost of Supply in Self-Sustaining Logistics Networks

Eva Regnier, Jay Simon, Daniel Nussbaum, Aruna Apte, and John Khawam
Naval Postgraduate School

Platform Design for Fleet-Level Efficiency: Application for Air Mobility Command (AMC)

Jung Hoon Choi, Parithi Govindaraju, Navindran Davendralingam, and William A. Crossley
Purdue University

Improving DoD Energy Efficiency: Combining MMOWGLI Social-Media Brainstorming With Lexical Link Analysis (LLA) to Strengthen the Defense Acquisition Process

Ying Zhao, Don Brutzman, and Douglas J. MacKinnon
Naval Postgraduate School

Addressing Counterfeit Parts in the DoD Supply Chain

Jacques S. Gansler, William Lucyshyn, and John Rigilano
University of Maryland

Wave Release Strategies to Improve Service in Order Fulfillment Systems

Erdem Çeven and Kevin Gue
Auburn University

Issues and Challenges in Self-Sustaining Response Supply Chains

Aruna Apte, John Khawam, Eva Regnier, Jay Simon, and Daniel Nussbaum
Naval Postgraduate School

Lead Time Demand Modeling in Continuous Review Supply Chain Models

Barry R. Cobb, *Virginia Military Institute*
Alan W. Johnson, *Air Force Institute of Technology*

Improving Multi-Component Maintenance Acquisition With a Greedy Heuristic Local Algorithm

Sifat Kalam and Kash Barker, *University of Oklahoma*
Jose Emmanuel Ramirez-Marquez, *Stevens Institute of Technology*

An Internal, Demand-Side Approach Toward Implementing Strategic Sourcing:



Political, Legal, and Economic Considerations

John Fallon, *University of Maryland, University College*
Timothy Reed, *Beyond Optimal Strategic Solutions*

Optimizing Causes of Procurement Cost Through Strategic Sourcing: The Impact of Rate, Process, and Demand

Timothy Reed, *Beyond Optimal Strategic Solutions*
Michael E. Knipper, *United States Air Force*
John Fallon, *University of Maryland, University College*



Wave Release Strategies to Improve Service in Order Fulfillment Systems¹

Erdem Çeven—Çeven is a PhD candidate in the Department of Industrial & Systems Engineering at Auburn University. His primary research goals are directed toward controlling logistics systems, especially in order fulfillment systems.

Kevin Gue—Gue is the Tim Cook Associate Professor in the Department of Industrial & Systems Engineering at Auburn University. He graduated from the U.S. Naval Academy in 1985 with a bachelor's degree in mathematics. He received his PhD from the School of Industrial & Systems Engineering at Georgia Tech in 1995. From 1995 to 2004, he was on the faculty of the Graduate School of Business & Public Policy at the Naval Postgraduate School. Dr. Gue's research interests include logistics modeling and optimization, with applications in distribution, warehousing, and material handling. He is a past president of the College-Industry Council on Material Handling Education.

Abstract

Using the Defense Logistics Agency's current service performance metric Next Scheduled Departure, we develop methodologies for establishing the optimal timing of order releases in a distribution center so that customers receive supplies sooner. We present a simulation model to test these methodologies and to show that setting wave release times accordingly can significantly improve service performance for systems subject to stationary and non-stationary arrivals.

Introduction

Continuing fiscal struggles in the federal government have made “do more with less” the operating mode of almost every Department of Defense (DoD) organization. The Defense Logistics Agency (DLA) and its distribution centers are no exception. In DLA's case, there is increasing pressure to provide superior service with the same or fewer labor resources. By “superior service,” we mean rapid response to customer requisitions.

“Operational availability” (A_o) of a system is defined as the fraction of time or probability that a system's capabilities will be available for operational use (“Operational Availability Handbook,” 2003). A_o is a function of “uptime” and “downtime,” the latter being mainly determined by Mean Logistics Delay Time (MLDT). Intuitively, reducing the flow time—the time between arrival of the order and the time it is ready to ship—decreases MLDT and therefore improves A_o . In general, the more quickly the logistics system responds to the requests, the higher the availability of end items because total downtime is reduced through reduced MLDT.

If a part is in stock, logistics delay time is comprised of two main components: warehouse processing time and transportation time. These two processes (warehousing and transportation) meet at the shipping dock. Doerr and Gue (2011) observed that warehouse operations are effectively “continuous,” in that completed orders arrive at a shipping dock, more or less, in a continuous stream. By contrast, transportation is a cyclical process, due to the need to achieve economies of scale. Thus, for example, package carriers such as UPS and FedEx have a “nightly sort” and “next day” deliveries. Less-than-truckload (LTL) carriers, which transport larger shipments, also operate according to a daily, cyclical model.

¹ This research has been funded under NPS-BAA-11-002 at Acquisition Research Program at the Naval Postgraduate School. Award Period: November 27, 2011–November 26, 2012.



To coordinate the internal, continuous operations of its DCs with the cyclical transportation schedules of its transportation providers, DLA uses a metric called Next Scheduled Departure (NSD), which measures the fraction of orders arriving during a specified 24-hour period that are processed before a specific truck departure (Doerr & Gue, 2011). By definition, an increase in the metric means more orders make their last departing trucks, and that some customers receive their orders before they otherwise would have. For repair parts and mission-critical consumables, therefore, increasing NSD reduces MLDT and increases operational readiness.

One would think that making an order available to pickers as soon as it arrives would increase its chance of making it onto the next departing truck, but such a view misses the economies of scale in a picking operation. A worker picking a large batch is much more efficient than a worker picking a single order, and therefore his capacity is higher. Higher capacity reduces waiting for arriving orders, and therefore tends toward lower total sojourn times. The benefits of large batches, however, must be weighed against the queuing time necessary to form the batch. To strike this balance, DCs at DLA release orders in large batches called waves.

Despite the ubiquity of wave operations in commercial (and military) warehouses, there are no analytical models to determine the optimal number and timing of these waves, especially to maximize performance against deadline-oriented metrics such as NSD, which is used at DLA. (A thorough review of literature is given in Çeven and Gue, 2013.) The goal of our work is to improve the service performance and thus MLDT of a distribution center (operated by DLA, the Services, or a third party) by properly setting wave release times.

In the next section, we discuss fulfillment operations at DLA Distribution Center, Susquehanna, PA, a fulfillment center operated by DLA Distribution. We analyze its order flow data and the current wave release strategy. In the section Optimal Wave Release Policies, we introduce our approximation models and use those models to verify a simulation model. We discuss the details of the simulation model in the Simulation Study section and summarize our findings in our conclusion.

Wave Operations at DLA Distribution

DLA Distribution Susquehanna, Pennsylvania (hereafter, DDSP) is an extremely complex distribution operation handling more than one million stock-keeping units (SKUs) stored among dozens of warehouses. One of its main service offerings is Dedicated Truck, in which a customer requests specific times of delivery each day or week. Delivery times could be, for example, daily at 1600, or every Tuesday and Friday at 1200, depending on the customer. DDSP then establishes departure times for trucks leaving to each Dedicated Truck customer.

Because an order arriving just before truck departure cannot possibly be processed in time, DDSP establishes an internal cutoff time (or set of cutoff times, as appropriate) for each Dedicated Truck customer. Orders arriving before that time are due on the appropriate “next truck” and must be processed before it departs. Distribution centers in the DLA measure their service performance with NSD, which measures the percentage of orders arriving between consecutive cutoff times that make it on the assigned truck. Figure 1 illustrates how the metric works.






Figure 1. Orders Arriving Between Consecutive Cutoff Times Are Due on the Next Deadline

Managing the release of work to the system in such an environment is a difficult task, to say the least, and especially so in the presence of waves. In a typical distribution operation, including at DDSP, there are 2–6 waves per day, depending on the workload and number of destinations that must be served. Figure 2b shows the scheduled wave release times at DDSP at the time of our study—0000, 0400, 0930, and 1600. In addition to these scheduled releases, orders were occasionally released manually at around 0700 and 0900 to balance the workload.




Figure 2. Number of Orders Arrived and Released Within a Day

It is our experience that the number and timing of order releases is based on intuition and experience of management. Could NSD be improved if the release times were changed? What level of benefit is possible? Before looking at the details of DDSP, we present an overview of mathematical models to establish order release times.

Optimal Wave Release Policies

We first discuss some major results from Çeven and Gue (2013), in which the arrival process is assumed to be stationary with rate λ orders per unit time. The authors propose a fluid approximation model in which individual orders are indistinguishable. (They also specify in which conditions this approximation is valid.) To maintain stability, the server’s capacity is assumed to be $\mu > \lambda$.

Arriving orders accumulate in a Warehouse Management System (WMS) virtual queue until the next wave is released, at which time the quantity of orders in that wave decreases at rate μ until the wave is complete. Waves in this model are not allowed to overlap; that is, the current wave must be complete before a new wave can begin. While the server is working a wave, orders in the next wave accumulate, and the cycle continues. The

cutoff time for accepting orders is assumed to be equal to the truck departure time, which is a worst case in terms of NSD. This assumption is not necessary, but it makes the presentation much easier. Çeven and Gue (2013) discuss how to assign a realistic cutoff time.

NSD in a single-wave system is a function of the wave release time w_1 , the arrival rate λ , and the server capacity μ . By definition,

$$\text{NSD} = \frac{\# \text{ orders worked today that arrived today}}{\# \text{ orders that arrived today}} \quad (1)$$

When there is a single class of orders (which is a simplest version of DDSP's original problem), the server should finish the wave exactly at the deadline, and therefore $w_1 = 1 - \rho$. Thus, the optimal NSD for a single wave system is $w_1^* = \text{NSD}^* = 1 - \rho$. This proposition alone provides important insights. First, the server should begin the work as late as possible in order to allow as many orders as possible to make it into the wave. Second, the wave should finish exactly at the deadline. Furthermore, the result suggests that the optimal cutoff time equals the optimal wave release time $w_1 = 1 - \rho$, for which NSD would be 100%. This can also be argued intuitively: releasing the wave before the cutoff time means some orders arrive after the release but before the cutoff. These orders will certainly miss the truck. Releasing the wave after the cutoff time means some orders are in the wave but are not due on the next truck, which reduces system capacity for the orders that need to be processed immediately. Neither condition should be optimal, as the result shows mathematically. Çeven and Gue (2013) also address the multiple wave systems and determine the closed form optimal wave release times for a system with stationary arrivals:

$$w_j = \begin{cases} \frac{j-1}{N} & , \text{ for } \lambda = \mu \\ 1 - \frac{\rho^j - \rho^{N+1}}{1 - \rho^N} & , \text{ for } \lambda < \mu. \end{cases} \quad (2)$$

The authors observe that the first release time does not change, but later wave times adjust as the number of waves increases. As expected $w_N = 1 - \rho$ when $N = 1$, and w_N (and NSD) converges to 1 as the number of waves $N \rightarrow \infty$. That is, NSD improves as more waves are released, especially when expected utilization is high. As utilization increases, the equation suggests that the maximum possible NSD decreases, converging to $(N - 1)/N$. The models presented in Çeven and Gue (2013) can also be extended to reflect uncertainty in both daily workload and capacity uncertainty as well as daily non-stationary arrivals. Although their results provide insight into the importance of setting proper wave release times, they only partially address the problem faced by DDSP. This is because Dedicated Truck operations are only a portion of each day's workload at DDSP, so it is impossible to assess capacity devoted to these orders. Another reason is the fact that DDSP often receives orders days in advance of when they are scheduled to ship, and many orders remain in queue until near their deadline. Another complication that is not addressed by Çeven and Gue (2013) is the existence of multiple deadlines. Nevertheless, the results in Çeven and Gue (2013) provide us with the ability to simulate and test different wave release policies.

Simulation Study

Using a data set from DDSP, we analyzed the existing order arrival stream and wave release policy in order to generate input for the analytical models. Distribution centers of DLA typically have outbound processes that include picking, packing, order consolidation, and shipping. Example flow timing data is given in Table 1.



Table 1. Order Flow Timing Data Sample

OD ORD	SD DT	SD TI	MIT DT	MIT TI	LS DT	LS TI	PICK DT	PICK TI	PACK DT	PACK TI	OFFER DT	OFFER TI	MISSION DT	MISSION TI
8002R005	010061	35900	010060	21206	010060	22633	010060	35404	010060	53610	010060	61740	010060	01700
8002R007	010061	35900	010060	21207	010060	22633	010060	35052	010060	53737	010060	61740	010060	01700

The first entry in Table 1 refers to the order ID. All date and time fields are defined as Julian day and military time, respectively (e.g., 2010257 refers to September 17, 2010; 230140 refers to time 2301 and 40 seconds). The following two fields refer to the scheduled departure time followed by the arrival date and time of the order. The field RLS refers to the date and time that the order is released for picking. Pick completion date and time is given in PICK DT and PICK TI. The completion date and time of packing is given in fields PACK DT and PACK TI. Because orders wait for consolidation, there is a consolidation date and time stamp (given with OFFER DT and OFFER TI). The last two data fields correspond to the actual shipment date and time. Figure 3 is an illustration of order flow.




Figure 3. Timeline of an Order Through Arrival-to-Ship Process

We were provided with three months of order flow data from January 2010 to March 2010 for Dedicated Truck operations (DTK)—a total of 402,406 orders. Of those orders, 351,866 arrived in 2010 (87.44% of total) and 351,530 (87.36% of total) orders were shipped during the three-month interval and were the subject of our analysis.

The managers of DDSP reported that the overall system performance in NSD was around 72% over the three-month period (75.0% in January, 57.9% in February, and 70.4% in March 2011). We observe that NSD is highly variable throughout the length of study, within a range of [22.7%, 100%]. On average, 72% of the customer orders were fulfilled by their deadline; however, on some days NSD dropped below 60% (see Figure 4).

Before describing the simulation, we must cover one last detail. Figure 2 shows that the arrival stream to DDSP is highly non-stationary. Çeven and Gue (2013) show how to modify the basic wave release model to handle this case. The “discretized” version of their model assumes the arrival rate in a discrete period of time (in this case, 15 minutes) is stationary, but that the mean rate may change from hour to hour.

Below we discuss both intuitive and optimal wave release policies and show how proper release times can improve NSD. The simulation study serves two purposes: (1) to verify the analytical models of Çeven and Gue (2013), and (2) to demonstrate that service performance can be improved with proper wave release times.




Figure 4. Recorded Daily NSD and DDSP

We model the order fulfillment system as a three-stage queuing system corresponding to the picking, packing, and shipping processes. We assume 20 servers per stage and identical exponential processing time distributions in each stage. This choice is arbitrary, of course, but in the absence of real data (DLA does not collect processing time data), we had no justification for another choice. Arriving orders are stored in a virtual queue and released in the next wave. Once an order is released for picking, available workers start picking orders. Completed orders are sent directly to packing and then to shipping. Because daily workload at DDSP varies, we test different levels of utilization $\rho = 0.5, 0.75, 0.95$. We adjust the (exponential) processing rate to maintain the appropriate utilization. We assume four waves per day, as in the operations at DDSP at the time of the study.

Before applying different wave release policies, we verify the simulation model by comparing simulated NSD with NSD according to the analytical models. Using a stationary arrival stream, we determine the optimal release times for a single class, four-wave system for each utilization level. Optimal release times suggest NSD would be 96.7%, 88.6%, and 78.1% for $\rho = 0.5, 0.75, 0.95$. We insert the release times into the SIMIO simulation software and run the model for 30 simulated days, with three days of warmup and 100 replications. Figure 5 shows the results. The analytical model approximates the corresponding system's NSD within 1%.




Figure 5. Verification of the Simulation Model

Note. The red line indicates the approximated NSD by the analytical model. Black, blue, and green data points correspond to the average simulated NSD for $\rho = 0.5, 0.75, 0.95$, respectively.



Consistent with the results in Çeven and Gue (2013), average NSD drops as utilization of the system increases. We also observe that the variability in NSD for both policies increases as ρ increases.

Next, we consider a non-stationary arrival stream representative of the DDSP data, but we scale the arrival rates to achieve an appropriate utilization. We first test an intuitive policy in which each wave has the same wave length. Because the system will be busy $\rho = 1 - w_1$ of the time, an equal time policy divides this interval into four equal waves.

To test the analytical model, we use the same non-stationary arrival data and determine optimal release times and NSD*. We insert the optimal release times into the simulation model and estimate the NSD_e. Table 2 shows the release times for the optimal and equal time policies.

Table 2 shows the approximated NSD* of the analytical model. Similar to our results for stationary arrivals, the simulation results are close to the approximations (e.g., the approximation overestimates the NSD by around 3%). The optimal policy performs 9.6%, 5.9%, and 1.2% better than the intuitive equal time policy for different levels of utilization. Recall that the model suggested more evenly distributed releases as utilization increases. We observe this situation especially for $\rho = 0.95$.

Table 2. Simulation Results for Non-Stationary Arrivals

$\rho = 0.5$	w_1	w_2	w_3	w_4	NSD* (%)	NSD _e (%)
Optimal policy	12:00	18:05	21:18	22:53	96.4	93.3
Equal time policy	12:00	15:00	18:00	21:00	-	83.7
$\rho = 0.75$	w_1	w_2	w_3	w_4	NSD* (%)	NSD _e (%)
Optimal policy	06:00	12:36	17:45	21:18	87.3	85.2
Equal time policy	06:00	10:30	15:00	19:30	-	79.3
$\rho = 0.95$	w_1	w_2	w_3	w_4	NSD* (%)	NSD _e (%)
Optimal policy	01:12	07:30	12:47	18:24	78.3	75.9
Equal time policy	01:12	06:54	12:36	18:18	-	74.7

Conclusions

In this study, we have addressed order release problems in order fulfillment systems and shown that setting wave release times properly can improve NSD, and thus the operational availability of supported end items. In order fulfillment systems such as those operated by DLA, order releases should be timed to accommodate daily deadlines.

In a simulation study, we verified the analytical results in Çeven and Gue (2013) with both stationary and non-stationary arrival streams. We implemented those models to test optimal policies against an intuitive policy and showed that releasing waves optimally improves NSD. Although the complexity of operations at DDSP made direct analysis prohibitive, our results suggest that NSD could be improved with further investigation into the number and timing of order releases.

References

- Çeven, E., & Gue, K. (2013). *Optimal order release strategies and cutoff times for order fulfillment systems with deadlines* (Working Paper).
- Doerr, K., & Gue, K. R. (2011). A performance metric and goal setting procedure for order fulfillment operations. Forthcoming in *Production and Operations Management*.



Operational availability handbook: A practical guide for military systems, sub-systems, and equipments. (2003).

Acknowledgements

The authors thank managers and analysts at DLA Distribution, Susquehanna, PA, for providing data and for many helpful discussions about wave planning systems.





ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net