

DTIC® has determined on 9/12/2009 that this Technical Document has the Distribution Statement checked below. The current distribution for this document can be found in the DTIC® Technical Report Database.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited. *Per ONR*

© COPYRIGHTED; U.S. Government or Federal Rights License. All other rights and uses except those permitted by copyright law are reserved by the copyright owner.

DISTRIBUTION STATEMENT B. Distribution authorized to U.S. Government agencies only (fill in reason) (date of determination). Other requests for this document shall be referred to (insert controlling DoD office)

DISTRIBUTION STATEMENT C. Distribution authorized to U.S. Government Agencies and their contractors (fill in reason) (date of determination). Other requests for this document shall be referred to (insert controlling DoD office)

DISTRIBUTION STATEMENT D. Distribution authorized to the Department of Defense and U.S. DoD contractors only (fill in reason) (date of determination). Other requests shall be referred to (insert controlling DoD office).

DISTRIBUTION STATEMENT E. Distribution authorized to DoD Components only (fill in reason) (date of determination). Other requests shall be referred to (insert controlling DoD office).

DISTRIBUTION STATEMENT F. Further dissemination only as directed by (inserting controlling DoD office) (date of determination) or higher DoD authority.

Distribution Statement F is also used when a document does not contain a distribution statement and no distribution statement can be determined.

DISTRIBUTION STATEMENT X. Distribution authorized to U.S. Government Agencies and private individuals or enterprises eligible to obtain export-controlled technical data in accordance with DoDD 5230.25; (date of determination). DoD Controlling Office is (insert controlling DoD office).

NOTE: At the sponsor's request, the final report (Surge Operations for Distribution Centers) was included in the proposal for continuation of funding (attached).

Technical Proposal*— ONR BAA 08-001
Submitted to the Office of Naval Research

DYNAMIC RESOURCE ALLOCATION TO
IMPROVE SERVICE PERFORMANCE IN
ORDER FULFILLMENT SYSTEMS

Kevin R. Gue
Department of Industrial & Systems Engineering
Auburn University
Auburn, AL 36849
kevin.gue@auburn.edu
334-844-1425

Award Period: January 1, 2009 – December 31, 2010

***Institutional Representative:** Dr. John M. Mason, Associate Provost and Vice President for Research, 310 Samford Hall, Auburn University, 36849, Tel: 334-844-4438, Fax: 334-844-5953.

Dynamic Resource Allocation to Improve Service Performance in Order Fulfillment Systems

Abstract

We have recently developed methods for approximating the sojourn time distribution for customers or jobs entering a multi-server queueing network with general interarrival and processing times. Here we propose to continue the development of such models, and use them to allocate workers dynamically in an order fulfillment system to improve its service performance. This work is part of a long-term effort to develop a resource control system for order fulfillment systems, such as those found in distribution depots of the Defense Logistics Agency.

1 Order fulfillment systems

America's ongoing transition from a manufacturing- to a service-based economy has motivated much interest in service systems, both from a management and a design perspective. On the management side, there has been renewed interest in marketing and consumer behavior and in measuring customer satisfaction. On the design side, there has been much work in supply chain network design, supply chain coordination, and inventory management.

The focus of our recent and proposed research is order fulfillment systems, which are manufacturing or distribution operations designed to respond to deadline-oriented customer requests. The importance of these systems to the U.S. economy should be obvious: nearly every industrial or consumer product sold in our country passes through such a system on its way to market.

The work we are proposing is rooted in four levels of customer awareness illustrated in Figure 1. The most basic level acknowledges the customer's existence, but not much more. We might think of this as the old-school mentality that, "This would be a great company, if it weren't for the customers." The second level is Measurement, in which a firm begins to assess its performance with respect to the perception of its customers. For example, it might survey its customers to determine if they have been satisfied with respect to quality or on-time delivery. The third level, which we believe is where most competitive firms exist, is Management. Here, the firm establishes customer-oriented metrics and manages its existing operations with respect to them. The fourth level, which is the subject of our work, is the Design level, in which the firm transforms (designs) its internal operations to improve these customer-oriented metrics. Because transformation involves change (by definition) and change is hard, we contend few companies exist at this highest level.¹

So, exactly what does all this mean in practice? The Defense Distribution Center (DDC), which is the distribution arm of the Defense Logistics Agency, provides an excellent case study of the transition from level two to level three in customer awareness. In the 1990's, DDC began more precisely measuring the time to respond to customer orders as a way of assessing whether or not it was providing "good service" to its customers. The new "average cycle time" metric was used by senior commanders at DDC to assess the performance of its many distribution depots. In the early 2000's, DDC noted problems with the effects of

¹We plead some literary license with these assertions. To say such things scientifically would require much interesting research, which is not in the author's area of expertise.




Figure 1: Four levels of customer awareness.

the metric on internal depot operations, and proposed a new metric called *Next Scheduled Departure*, which more accurately reflected service performance to the customer. This metric is still in use today, and depot commanders use it to manage daily operations. Note, then, the transition from simple performance *measurement* (average cycle time of internal operations) to *management* of operations to improve a customer-oriented metric like NSD. DDC successfully moved into level three of customer awareness.

How then might DDC move to the next level, in which internal operations are transformed to improve customer-oriented metrics? This is the subject of our proposed research.

2 A theory of order fulfillment

At the most basic level, order fulfillment involves only two parties: Customer and Provider. Customer makes a request for a product or service with an implicit or explicit deadline, and Provider provides the goods in the expected condition by the expected time. The three basic functions are Order→Process→Deliver. The problem with this model is that it is extremely inefficient for most operations. For example, if upon every new order, a warehouse worker walks into the warehouse to pick a single item, and a truck leaves the dock with that single order to a single customer, we would have a very expensive order fulfillment system indeed.

A more efficient system uses economies of scale at two points: orders are batched before processing, which reduces processing costs, and processed orders are batched for delivery, which reduces transportation costs. The process is Order→Batch→Process→Batch→Deliver. In the warehousing literature, the effects of batching on order picking processes is well-researched and well-understood

(van den Berg and Gademann, 1999). Because orders are assumed to be exogenous to the system, there has been no research (that we know of) on how order batching affects the order stream itself (and we doubt that it does). To gain further economies of scale, a second batching operation occurs before transportation, which reduces delivery costs. Assigning orders to trucks is another name for vehicle routing problems, which are also well-researched and fairly-well understood.

But how, we ask, should the batching process which is transportation affect the internal processes of the warehouse? This is not well-understood, and, to our knowledge, has not been addressed in the literature.

3 Purpose

The goal of our research is to design dynamic resource allocation processes that improve customer service in order fulfillment systems. To be more specific, we seek dynamic worker allocation policies in a warehouse environment that get more orders on outbound trucks sooner.

We propose to investigate policies that require the sojourn time distribution for an order. Our research will make two contributions:

1. We will introduce to the academic and practicing communities dynamic worker allocation policies designed to improve *service performance* instead of throughput.
2. We will extend the theoretical development of approximation models for sojourn time distributions for multi-server queueing networks with general interarrival and processing times. We have recently made some significant advances in this area; our new work proposes to go further.

4 Progress to date

The general theme of the research is one of two we have been addressing under our current grant. In the course of our research in these two years, this particular subject (dynamic worker allocation for order fulfillment systems) has grown significantly and has led to two major theoretical advances, which we describe below.

To begin, we should say that the Next Scheduled Departure metric itself is new to the academic community, at least in the context of order fulfillment and

manufacturing systems. Manufacturing research especially has almost universally sought to optimize the classic measures of system performance, such as work-in-process inventory, throughput, or cycle time. In the warehousing literature, the objective is almost always throughput. Our work has been different, and we believe valuable, because it supposes a different, more customer-focused objective. By definition, Next Scheduled Departure is the percentage of orders arriving between cut-off times on two consecutive days that are shipped on a departing truck after the cut-off time on the second day. For example, suppose the order cut-off time is 1400 (see Doerr and Gue, 2006, for a discussion of how to set the cut-off time), and 1,000 orders arrive between 1400 on Day 1 and 1400 on Day 2. If 900 of those orders ship on the 1700 truck on Day 2 (the rest being still in process), then NSD is 90 percent for that day. As we have said, this metric is in use at DLA today.

The idea behind dynamic worker allocation came during some research with the Defense Distribution Depot in San Diego (DDDC) several years ago, when we noticed that about 5 percent of orders were ready to be shipped in the 30 minutes *after* the truck left everyday. (This insight led to the development of NSD, in fact.) The research question we posed at the time was, "Could internal operations at DDDC have been reorganized just before the last truck departed for the day, such that more orders got on that truck?" Notice that advancing that 5 percent of orders by 30 minutes would result in customers receiving them earlier by an entire day, due to the batch nature of transportation.

In our current grant, we began investigating just this question: Can workers be dynamically allocated in a warehouse to improve the NSD metric? Early simulation experiments, frankly, were inconclusive. We tried some naïve policies such as,

Policy 1 (*Naïve 1*) *At 1400 move 3 workers from Picking to Shipping, then move them back after truck departure,*

but this did not seem to improve performance consistently, and in any case not significantly. The problem was that the policy was "dumb": On some days, there was a significant queue in shipping at 1400, and the extra workers helped push through orders that otherwise would not have made the truck. But on other days, workers moved to shipping when there was no work for them to do, thus making them idle for a time and hurting performance in the long run.

These insights led us to formulate *state-dependent* policies, in which the state of the system (number of orders in the shipping queue, for example) determine when and how many workers to move. Naïve state-dependent policies, such as




Figure 2: A notional model of an order fulfillment system. Filled dots indicate workers.

Policy 2 (Naïve 2) *At 1400, if the shipping queue is more than 10 orders, move 3 workers from Picking to Shipping; otherwise, move none,*

showed enough promise in simulation to warrant continued investigation.

Eventually we found that the most promising policies require an understanding of the sojourn time distribution for an order. For example, we needed to know what was the probability that an order arriving at, say 1230, would finish before the truck left at 1700. If the probability is low, then we might “abandon” that order and shift resources to more promising orders. Answering this sort of question requires a distribution of sojourn time, not just an expected value. This has led us to two theoretical advances.

4.1 A steady-state sojourn time model

We model the warehouse as a serial line of three workstations (picking, packing, and shipping), with multiple workers for each workstation, or, more generally, as a network of multi-server queues. Our search of the literature revealed that there existed no methods for approximating the sojourn time distribution in such a network.

We have developed an approximation model for the sojourn time distribution in a network of multi-server queues, when the interarrival and processing times can take on general distributions. We believe we are the first to have developed such a model. Below we sketch the procedure and give a glimpse of the results. Details and mathematics are available in Gue and Kim (2008a).

Our method is based on characteristics of phase-type distributions, which allow us to approximate general distributions in a way that we can take convolutions of waiting and processing time distributions at each stage to arrive at a final sojourn time distribution. Our work builds on existing research in matrix-geometric

methods by Neuts (1981), Asmussen and Møller (2001), and You et al. (2002). Very briefly, the procedure is

1. Use the QNA method of Whitt (1983) to establish interarrival processes to all workstations.
2. Convert each $G/G/c$ representation of a workstation (for picking, packing, shipping) to a corresponding phase type model $Ph/Ph/c$, based on the squared coefficients of variation C_a^2 and C_s^2 .
3. Compute the initial probability vectors and infinitesimal generators (γ_i, R_i) for the waiting time distribution of each $Ph/Ph/c$ queue and compute corresponding vectors and generators (β_i, S_i) for processing time, according to Asmussen and Møller (2001). For the number of phases in each $Ph/Ph/c$ distribution, use the appropriate value of $\lceil 1/C^2 \rceil$.
4. Use the (β_i, S_i) and (γ_i, R_i) with the method of You et al. (2002) to generate an initial probability vector and infinitesimal generator of sojourn time in the system (ζ, Q) .
5. Solve $F(t) = P(T \leq t) = 1 - \zeta e^{(Q)t} \mathbf{e}$ to obtain the CDF of the sojourn time distribution.
6. Repeat Steps 3–5, this time with number of phases $\lfloor 1/C^2 \rfloor$ for each distribution. Let the resulting CDF be $G(t)$.
7. Compute the mixed CDF, $H(t) = \alpha G(t) + (1 - \alpha)F(t)$, where α is an interpolation coefficient, which we show how to determine in Gue and Kim (2008a).

Our method produces distributions very similar to those established in a discrete event simulation. For a prototypical system with 3 stages and 6 workers per workstation (see Figure 3), the results are quite close (Figure 4). We did extensive analysis with Anderson-Darling tests to show that the approximation model produces distributions statistically identical to those derived from simulation. A full description of the method and the results is in Gue and Kim (2008a).

With the steady state sojourn time distribution in hand, we made a significant discovery with respect to the performance metric NSD, which enables us to compute the expected NSD for an order fulfillment system directly from the sojourn




Figure 3: A system with 6 workers per workstation in 3 workstations.




Figure 4: The pdf and CDF of sojourn time for the example system in Figure 3. Graphs on top correspond to $E[T]$ values of Figure 3; graphs on the bottom to values corresponding to utilization $\rho = 0.5$.




Figure 5: Computing expected NSD from the steady state sojourn time distribution.

time distribution. In a coming paper, we will show that

$$NSD = \frac{1}{24} \int_{\delta}^{\delta+24} P[T < t] dt, \quad (1)$$

where $P[T < t]$ is the CDF of the sojourn time distribution, $0 \leq t \leq 24$, and δ is the amount of time between the cut-off time and the departure time. Figure 5 shows the relationship graphically.

The discovery is important because it allows managers, with a relatively simple calculation, to estimate performance on NSD for any cut-off time under consideration. Doerr and Gue (2006) discuss how to set the cut-off time with respect to motivating workers to achieve an NSD goal; our result makes this process much easier. Equation 1 will also help us with the exploration we propose in Section 5.

4.2 A state-dependent sojourn time model

Our second major advance has been a *state-dependent* model of sojourn time distributions, again, for multi-server queueing networks with general interarrival and processing time distributions. We believe ours is the first such model. By state-dependent, we mean the following. Suppose an order is twentieth in queue when

it arrives to a queueing network. Furthermore, suppose we know the number of orders in every queue in the system at the time of arrival. What is the sojourn time distribution for the job that just arrived?

As before, the model uses phase-type distributions and matrix-geometric methods. The procedure is as follows:

1. Approximate the first service time distribution as a corresponding phase-type distribution, (β, S) , based on the SCV, C_s^2 .
2. Approximate the first waiting time distribution as a corresponding phase-type distribution, (α_1, W_1) .
3. Approximate the second waiting time distribution as a corresponding phase-type distribution, (α_2, W_2) based on the first initial probability vector, α_1 .
4. Approximate the following waiting time distributions as appropriate phase-type distributions, (α_i, W_i) , $i \geq 3$, according to the same procedure.
5. Generate the initial probability vector and infinitesimal generator of sojourn time distribution of the system (γ, K) using the convolution property of the phase-type distribution.
6. Solve $F(t) = P(T \leq t) = 1 - \gamma e^{Kt} \mathbf{e}$, $(t \geq 0)$ to obtain the CDF of the sojourn time distribution.

A full description of the model and detailed results are in Gue and Kim (2008b).

With a model for state-dependent sojourn time distributions, we can compute *for any job, at any time* its probability of “success,” or the probability that it will make it on the target truck. This gives us a powerful tool for real-time decision-making in order-fulfillment systems.

5 Statement of Work

We propose to move forward in two main directions: (1) To explore the behavior and control of order fulfillment systems using policies based on sojourn time distributions, and (2) to extend our approximation models to account for non-stationary distributions commonly seen in the order streams of distribution centers, such as those found throughout the Defense Logistics Agency.

5.1 Literature review

The use of cross-trained workers to improve the performance of manufacturing systems has been investigated by several authors. Askin and Chen (2006) propose two types of worksharing policies—Dynamic assembly-Line Balancing (DLB) and Moving Worker Modules (MWM). MWMs have fewer workers than machines; DLB systems have the same number of machines and workers, and workers do not move. DLB has fixed tasks and shared tasks; the former are assigned to a designated worker while the latter can be carried out by either of an adjacent pair of workers. In addition to this classification, Hopp and Van Oyen (2004) suggest a “floating workers” category, in which a fixed worker occupies a machine and some cross-trained workers move dynamically to machines most needing them. Our review will cover only the work closely related to our proposed research. The floating worker category is closest to our problem, but different in that we can have many workers at any machine.

One policy of particular interest is the half full buffer (HFB) control policy for dynamic assembly-line balancing. When the buffer between two machines is more than half full, a flexible worker helps a downstream worker, or vice versa. Ostalaza et al. (1990) developed a Markov chain model for two stage systems and presented simulation results for longer lines. McClain et al. (2000) tested the HFB control policy to move flexible workers for shared tasks. Gel et al. (2002) proposed a more complicated HFB control policy, which considers general processing times and different worker speeds. Askin and Chen (2006) and Chen and Askin (2006) suggested another threshold heuristic rule—smallest R , no starvation (SRNS).

Floating worker systems have been examined by Sennott et al. (2006), who considered serial production lines in which each station has a specialist, and one fully cross-trained generalist to be assigned to any station. They developed Markov decision models including set-up costs, holding costs, and set-up time and showed that the flexible system performs better than a static system. Andradóttir et al. (2001) examined a tandem line with two stations and two cross-trained workers using a Markov decision model.

Our work is different than previous work in three important ways:

1. Existing work considers only the traditional system performance metrics of throughput, cycle time, and WIP. Our work is specifically directed toward the service-oriented metric NSD.
2. Almost all existing work assumes exponential processing times. The exception above is Gel et al. (2002), which admits general processing times.

3. All existing work assumes a serial production line, and most often just a two-stage system. Our models will allow general networks of reasonable size.

5.2 Worker allocation policies

We propose to investigate worker allocation policies both for distribution centers, which can be modeled as a serial line of multi-server queues, and for more general manufacturing and logistics networks, which may be modeled as more general networks of multi-server queues.

5.2.1 For distribution centers

We have already developed conceptual models for three classes of dynamic worker allocation policies, which we call the *flushing policies*, *cascade policies*, and *concurrency policies*.

The flushing policy is easy to implement, and is based on the following insight: Orders that “just miss” the truck are probably located near the shipping area shortly before the truck leaves, so it seems reasonable to check the queue in shipping a fixed time before the truck leaves and decide how many (if any) workers to move to Shipping from Picking. To decide how many workers to move, we look at the $P(\text{success})$ of the last order in the shipping queue. If it is below a threshold value, we move sufficient workers to Shipping such that its recalculated $P(\text{success})$ is above the threshold.

Policy 3 (*Flushing*) *At a fixed time each day, compute $P(\text{success})$ for the last order in the shipping queue. If it is below a threshold value, move sufficient workers to Shipping such that its recalculated $P(\text{success})$ is above the threshold. Restore the system to its original configuration upon truck departure.*

We have performed extensive simulation experiments for this policy, and have observed potential increases in NSD performance between 5 and 10 percent. In a service context, this means that 5–10% of customers would be receiving their orders one day earlier than they would be without dynamic worker allocation. We strongly suspect retail firms (and DLA) will be interested in these results.

But we have much to do, even with this simple policy. For example, the policy is controlled by three parameters (time at which to switch workers, the threshold probability, and the number of workers to switch), and we do not yet

understand how these parameters affect the results. Our simulation work has used only “reasonable values” for these parameters.

We propose to investigate the performance of the flushing policy for dynamically allocating workers in a serial queueing network, which we believe fairly represents a distribution center operation. Our investigation will be primarily simulation-based, due to the complexity of the system. Our objective is to understand the relationship between switching time, threshold probability, and the number of workers switched. We anticipate offering to the research and practicing communities a real-time decision-making system with the potential to improve order fulfillment times by an entire day for 5–10 percent or more of its customers.

We will also investigate a modified version of the Flushing policy, which we call Multi-Flush.

Policy 4 (*Multi-Flush*) *At multiple fixed times per day, execute the Flushing Policy above.*

The intuition here is that on some days a shift of workers at the normal switching time may not afford sufficient time to work down the queue in Shipping. Would it be better, then, to do a mid-day switch, for example, to work down the queue early on?

Such a policy may be especially useful when workload is not steady throughout the day, as is the case for most every order fulfillment system. We address the problem of sojourn time estimation under non-stationary arrival distributions below.

Policy 5 (*Cascade*) *At a fixed time each day, execute the flushing policy from Picking to Packing. At a later time, execute flushing from Packing to Shipping.*

The idea behind this policy is to “follow the work” that will just finish before departure, creating a sort of “bow wave” of orders working their way through. Parameters to modify here are the same as for the Flushing Policy, plus we must decide on a second time of execution.

Policy 6 (*Concurrence*) *At multiple fixed times per day, consider switching workers from Picking to Packing, and from Packing to Shipping, simultaneously.*

This class of policies has considerably more flexibility, but will also be more difficult to design and control. We believe that, as with the Multi-Flush and Cascade policies, Concurrence will be especially valuable for non-stationary arrival distributions.




Figure 6: A simple general network.

5.2.2 For manufacturing and logistics networks

So far we have addressed only serial processes, which are a good approximation of a typical warehouse. Manufacturing and logistics networks can also be modeled as networks of queues, and our methods might be used to establish control schemes to improve their performance. Therefore, there is a need to design resource allocation policies when the network is not a serial line. The general network control problem is more complicated because it is not clear how to move workers between workstations, nor are the implications of those moves clear. We propose to investigate worker allocation policies for simple general networks, such as the one in Figure 6, and move to more complicated networks after developing some insight.

5.3 Theoretical extensions

We also propose to extend our theoretical models in two ways. The first is to address non-stationary arrival distributions. Thus far, we have developed models only for the case of stationary arrival distributions. The arrival patterns to most service systems, however, are more complex. For example, data from a DLA depot with whom we worked indicated surges of orders around 0500 and 1400 each day, which created a bimodal distribution. We anticipate that we will extend

our existing models in another dimension of time, creating, for example, a CDF of the form $F(t_1, t_2)$, where t_1 is the arrival time (of day) and t_2 is the sojourn time. The result will be a sojourn time *surface*, with which we will construct worker allocation policies.

The second theoretical extension will incorporate non-zero switching times. Currently, our models assume that workers move from one workstation to another instantaneously. In practice, of course, repositioning workers creates idle time during the transition, and we must account for this.

6 Future Naval relevance

Our work on sojourn time distributions, although fairly technical, has a very practical application in real order fulfillment systems. We are trying to build the components of a resource control system for order fulfillment systems, which could be used in industry or the military.

To that end, we have had conversations with the Deputy Commander of the Defense Distribution Depot, Susquehanna, PA (DDSP), and he has expressed interest in developing such a control system. We have much to do before this would come to fruition:

1. We must accomplish the work in this proposal, which should more firmly establish the theoretical underpinnings of the control policies.
2. We must do extensive testing with data from DDSP. We have already received a data set from DDSP, and have plans to test our models with it in the next phase of the research.
3. We would need to build a software tool. Such a tool would be useful with even the most basic functions of tracking the work content within the distribution center and giving managers a way to identify bottlenecks and idleness.
4. Finally, integration and testing at the field site (DDSP) would be required.

To reiterate, we are not proposing to accomplish all of this in this proposal. Rather, we thought it would be helpful to point out that the current, proposed research is part of a larger effort to bring research and technology to practice in the DoD.

7 Vita

Kevin R. Gue is Associate Professor in the Department of Industrial & Systems Engineering at Auburn University. Prior to 2004, he was Associate Professor of Logistics in the Graduate School of Business & Public Policy at the Naval Postgraduate School. He graduated from the U.S. Naval Academy in 1985 and served as an officer in the submarine community for 5 years. After leaving active duty in 1990, he attended graduate school at Georgia Tech, where he received his Ph.D. in Industrial Engineering in 1995. His current research interests include facility logistics, warehousing, distribution, and the effects of performance metrics on design and operations. He has published articles in *Naval Research Logistics*, *Operations Research*, *Transportation Science*, *IIE Transactions*, and other journals. Previous research sponsors include the Office of Naval Research, the National Science Foundation, and the Defense Logistics Agency. Dr. Gue is currently Past-President of the College-Industry Council on Material Handling Education.

References

- Andradóttir, S., Ayhan, H., and Down, D. (2001). Server assignment policies for maximizing the steady-state throughput of finite queueing systems. *Management Science*, 47(10):1421–1439.
- Askin, R. and Chen, J. (2006). Dynamic task assignment for throughput maximization with worksharing. *European Journal of Operational Research*, 168:853–869.
- Asmussen, S. and Møller, J. R. (2001). Calculation of the Steady State Waiting Time Distribution in GI/PH/c and MAP/PH/c Queues. *Queueing Systems*, 37:9–29.
- Chen, J. and Askin, R. (2006). Throughput maximization in serial production lines with worksharing. *International Journal of Production Economics*, 99:88–101.
- Doerr, K. H. and Gue, K. R. (2006). A Performance Metric and Goal Setting Procedure for Order Fulfillment Operations. Working paper, Naval Postgraduate School, Monterey, CA.
- Gel, E., Hopp, W., and Van Oyen, M. (2002). Factors affecting opportunity

- of worksharing as a dynamic line balancing mechanism. *IIE Transactions*, 34:847–863.
- Gue, K. R. and Kim, H. H. (2008a). An Approximation Model for Sojourn Time Distributions in Multi-Server Queueing Networks. Working paper, Auburn University.
- Gue, K. R. and Kim, H. H. (2008b). An Approximation Model for State-Dependent Sojourn Time Distributions in Multi-Server Queueing Networks. Working paper, Auburn University.
- Hopp, W. and Van Oyen, M. (2004). *Agile workforce evaluation: a framework for cross-training and coordination*. *IIE Transactions*, 36:919–940.
- McClain, J., Schultz, K., and Thomas, L. (2000). Management of worksharing systems. *Manufacturing and Service Operations Management*, 2(1):49–67.
- Neuts, M. F. (1981). *Matrix-Geometric Solutions in Stochastic Models*. The Johns Hopkins University Press, Baltimore and London.
- Ostalaza, J., McClain, J., and Thomas, L. (1990). The use of dynamic (state-dependent) assembly-line balancing to improve throughput. *Journal of Manufacturing and Operations Management*, 3:105–133.
- Sennott, L., Van Oyen, M., and Iravani, S. (2006). Optimal dynamic assignment of a flexible worker on an open production line with specialists. *European Journal of Operational Research*, 170:541–566.
- van den Berg, J. P. and Gademann, A. J. R. M. N. (1999). A literature survey on planning and control of warehousing systems. *IIE Transactions on Scheduling & Logistics*, 31(8):751–762.
- Whitt, W. (1983). The Queuing Network Analyzer. *Bell Systems Technical Journal*, 62(9):2779–2815.
- You, J. U., Chung, M. Y., and Sung, D. K. (2002). An Approximation Method for Sojourn Time Distributions in Telecommunication Networks. *Computer Communications*, 25:1283–1296.



AUBURN UNIVERSITY
SAMUEL GINN COLLEGE OF ENGINEERING
Engineering Experiment Station

August 3, 2009

ATTENTION:

AGENCY: Defense Technical Information Center

ADDRESS: 8725 John J Kingman Rd. Ste 0944
Fort Belvoir, VA 22060-6218

CONTRACT/GRANT: N00014-07-1-0305

The following approved information is submitted for the closing of the subject contract/grant on behalf of Dr. Kevin R. Gue in the Department of Industrial & Systems Engineering:

Final report or (Copy of transmittal letter) w/SF298

New Technology Summary Report (NASA NTSR)

Final Invention Statement and Certification

DD Form 882/MSFC Form 4204

If you need any additional information, please give me a call at (334) 844-2297.

Joanne Tolbert
Financial Associate II