Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

A SIMULATION ANALYSIS OF THE VARI-METRIC REPAIRABLE INVENTORY OPTIMIZATION PROCEDURE FOR THE U.S. COAST GUARD

Michael B. Zamperini

Mathematics Department
27 Mohegan Avenue
U.S. Coast Guard Academy
New London, CT 06320, U.S.A .

Michael Freimer

Supply Chain and Information Systems Department
Smeal College of Business Administration
509N Business Administration Building
Penn State University
University Park, PA 16802, U.S.A.

ABSTRACT

This paper documents a simulation study undertaken to gain insights into the Vari-Metric multi-echelon repairable inventory optimization method. The method was analyzed in the context of the Coast Guard's fleet of fixed and rotary wing aircraft, for which operational availability is the key performance metric. Failure rates of parts in this system exhibit variance-to-mean ratios higher than one, and we outline a procedure for generating a failure arrival process described by a negative binomial distribution. Previous studies of the Vari-Metric model examined a single repairable part; in this study we analyzed a system comprised of three repairable parts. This allowed us to gain insights into how the Vari-Metric procedure selects part stock levels to attain a desired level of system availability. Further analysis of the simulation model allowed us to examine the efficient frontier for this multi-criteria problem (maximizing spare part availability rates while minimizing the cost of part inventories).

1 INTRODUCTION

This research was motivated by the U.S. Coast Guard's desire to evaluate the Vari-Metric multi-echelon repairable inventory optimization method put forth by Sherbrooke (1992). The Vari-Metric procedure calculates the expected backorders for a given inventory arrangement, then employs a greedy algorithm to iteratively add parts where they produce the greatest expected reduction in backorders per dollar invested until a target system availability level is obtained. The Coast Guard wished to apply this procedure to the spare parts inventories for its fleet of fixed and rotary wing aircraft located at different air stations throughout the country. We conducted a simulation study to evaluate the use of the Vari-Metric procedure in this context. We used a commercially available software implementation of Vari-Metric for this study. Our simulation model comprised 26 units of the same style of aircraft, with three repairable parts operating at five geographically separated stations, and with a central warehouse and part repair facility. A complete description of this model can be found in Zamperini (2005).

Sherbrooke's Vari-Metric work was also largely focused on aircraft spare part inventories, and his studies indicated that it was common for failure data from individual aircraft parts (number of failures per time period) to have variance to mean ratios (VMRs) greater than one. This suggested that a Poisson model may not adequately describe the variance in the failure arrival process. Following work commenced by Graves (1985), Sherbrooke improved on early versions of his Metric model by employing the negative binomial distribution to more accurately reflect the variance in part failure processes. Since the negative binomial distribution has two parameters, one can fit both the mean and variance of a unimodal data set, although the VMR must be greater than one. The commercial version of the Vari-Metric model employed by the Coast Guard assumes a negative binomial distribution for failures, and in fact the failure data from our study was well-described by this distribution. In this paper we document the method for generating an arrival process from a negative binomial distribution used in our simulation, and we share some of the insights obtained through our simulation experiments.

2 THE SIMULATION MODEL AND FAILURE PROCESS

We constructed a model of the Coast Guard's inventory system using the Arena® simulation package.

2.1 Modeling Assumptions

The model tracked the operability of three parts on each of 26 aircraft. Upon failure of a part, we assumed that the aircraft was completely inoperable and remained in that state until the appropriate part became available for installation. The failed part was immediately sent to the repair depot for repair and subsequent reentry into the inventory system. Upon completion of its repair, a part was shipped to an air station if any were below their target inventory levels. Otherwise the part was held at the warehouse until such an inventory shortfall arose. Our simulation model could also be set up to allow lateral transfers of parts from base to base. Lateral transfer was triggered if one base was out of a part, the warehouse did not have that part ready to issue, and another base did have the part available.

We tracked shipping and repair times, but we assumed a part installation time of zero once the appropriate part was available. Our key performance metric was *operational availability* (Ao). An aircraft was considered operationally available if it was not down due to a lack of spare parts. Ao was measured as a fraction of the overall number of aircraft: one minus the time-average fraction of aircraft waiting for a spare part. We tracked the system Ao as well as availability by location and individual part.

2.2 The Failure Process

In addition to the work of Sherbrooke and Graves, other research has been conducted on the topic of arrival processes exhibiting extra-Poisson variability. Sokhan-Sanj et al. (1999) proposed the use of a hyperexponential distribution to simulate highly variable part movements in a semiconductor fabrication setting. By accurately capturing the actual system variability, they were able to eliminate an undesirable safety factor that had been commonly used in previous simulation studies. Pichitlamken et al. (2003) studied the variability of incoming calls to a call center per half hour period. They described the potential use of the negative binomial distribution but selected a method employing a gamma distribution and inducing correlation between time periods. This method significantly outperformed a Poisson process with a set of time-dependent arrival rate parameters when compared with known system characteristics.

The Vari-Metric model can also incorporate extra-Poisson variability in the arrival process. Although its basic assumption is that demand is stationary Poisson, this can be generalized to allow non-stationary demand with an arrival rate that drifts over time. Sherbrooke (1992) suggests that the negative binomial distribution may be an appropriate representation of demand, however the probabilistic structure of the demand process is not made explicit in the generalized Vari-Metric model. Instead the generalized model approximates the steady-state distribution of the number of outstanding repairs with a negative binomial distribution having a VMR matching that of the demand process. Our analysis employed the generalized model.

We examined part failure data for several Coast Guard aviation parts. These data described the number of failures per week over a multi-year period for each part. Matching Sherbrooke's findings, each of these parts had VMRs greater than one, ranging from 1.16 to 2.07. The number of failures per week was well described by a negative binomial distribution, and the data sets exhibited no appreciable trends or serial correlation. We therefore created part failure processes for our simulation in which the number of failures each week was negative binomial with mean and variance matching the observed parameters. This was as close as we could come to matching the probabilistic assumptions of the generalized Vari-Metric model.

We used the following fact to create the arrival processes: if θ is generated from a gamma distribution with parameters k and $p/(1-p)$, and if X is generated from a Poisson distribution with parameter θ, then X is negative binomial with parameters k and p (Bratley et al. 1987). Let $\overline{X}_i$ and S_i^2 be the empirical mean and variance of the weekly number of failures for the i^{th} part type. Set $k_i = \left(\overline{X}_i\right)^2 / \left(S_i^2 - \overline{X}_i\right)$ and $p_i = \overline{X}_i / S_i^2$. During the j^{th} week of the simulation, draw θ_{ij} from a gamma distribution with parameters k_i and p_i. Failures of the i^{th} part during the j^{th} week are then generated via a Poisson process with rate parameter θ_{ij}, and the number of failures of part i during week j is distributed negative binomial with parameters k_i and p_i (i.e. with mean $\overline{X}_i$ and variance S_i^2).

We generated interarrival times for failures of part i during week j from an exponential distribution with mean $1/\theta_{ij}$. Since θ_{ij} was only used to draw interarrival times within week j, if any random draw resulted in an arrival time after the end of week j, we made another draw for the beginning of week j+1 using $1/\theta_{i,j+1}$. If the resulting draw was greater than one week, we made yet another draw using $1/\theta_{i,j+2}$. We continued in this way until a random draw of less than one week was obtained, say in week $j+k$. Then the final interarrival time was the time remaining in week j, plus k-1 weeks and the value of the random draw for week $j+k$. The purpose of this procedure was to ensure that the distribution of the number of failures in any week j conditioned on θ_{ij} was Poisson with parameter θ_{ij}, so the unconditioned distribution was negative binomial with parameters k_i and p_i. The procedure also reflects Sherbrooke's notion of a failure process that is locally Poisson but whose arrival rate drifts randomly. We tested the effect on our model of using this arrival process as compared with a Poisson process; the results are described in Section 3.4.

The method described in the preceding paragraphs is an unusual procedure for generating non-stationary arrivals. A more common approach is to generate arrivals from a non-homogeneous Poisson process via thinning (Law and Kelton 2000). However our purpose was to ensure that the number of arrivals each week had a negative binomial distribution, which is not the case for a non-homogeneous Poisson process.

3 SIMULATION RESULTS

In this section we summarize our experiments with the Vari-Metric software and our Arena® simulation model. Coast Guard analysts had expressed concern about the appropriateness of Vari-Metric recommendations under several scenarios. Based on these concerns, the first set of experiments used Vari-Metric to generate spare part stock policies for a target Ao under a variety of conditions, then checked the performance of these policies using our simulation model. (A policy consists of target stock levels for each of the three parts at the five stations and the central warehouse.) The results of this analysis led us to a second experiment in which we used the simulation model to examine the efficient frontier with respect to inventory cost and target Ao. A third set of experiments was designed to test the effects of three assumptions made by the Vari-Metric procedure: negative binomial failures, infinite repair capacity, and deterministic repair times. Finally, we used the simulation to investigate the interaction of failure VMR and part price. By examining how individual parts contributed to the overall system Ao, we gained valuable insight which led to a new method for dealing with high VMR parts in the Vari-Metric procedure.

3.1 Exercising Vari-Metric

We investigated the performance of the Vari-Metric recommended policies under eight scenarios generated from three variables with two settings each. We chose part failure VMR, part price, and part failure rate as the three control variables based on concerns expressed by the Coast Guard analysts. While it makes sense to maintain larger stock levels for parts with high variability and to provide cheap insurance by buying large numbers of the less expensive parts, some analysts hypothesized that Vari-Metric "overbought" high VMR and low price parts. Failure rate was chosen as the third control variable as the typical repairable part is slow moving as well as expensive.

We increased the VMR and failure rate from their nominal levels and decreased part price, so the -+- case reflected our actual part data. As described in the previous section, the simulation model tracked the availability of three part types for each aircraft. In these experiments, each control variable was altered for a single part type only. VMR was changed for part 1 only, while price was altered for part 3, and failure rate was changed for part 2. The high and low levels for each control variable and their associated parts are shown in Table 1.

Table 1: Control Variables

Variable	Level	Symbol	Part Number / Value
VMR	High	+	Part 1 VMR = 3
	Nominal	–	Part 1 VMR = 1.45
Price	Nominal	+	Part 3 price = $158,716
	Low	–	Part 3 price = $150
Failure Rate	High	+	Part 2 failure rate = 12 per week
	Nominal	–	Part 2 failure rate = 0.156 per week

In each of the eight scenarios we used Vari-Metric to generate a stock policy with a target Ao of 85%. (We also performed an analysis in which the target policy was 95%; the data are not shown here.) Table 2 reports the Ao predicted by the Vari-Metric procedure, as well as the steady-state Ao estimated with the simulation model. For each scenario we performed 100 independent simulation replications of 200,000 days with a warm-up period of 20,000 days.

Our analysis indicates that under the assumptions of our failure process, the Vari-Metric recommendations may be conservative. In each scenario the Ao predicted by Vari-Metric was lower than the value estimated by the simulation; in the base case the difference was more than five percentage points. (We observed a similar effect when the target Ao was 95%, although the deviations were smaller.) We speculate that the reason for this discrepancy lies with Vari-Metric's assumption that the number of outstanding repairs has a negative binomial distribution with a VMR matching the arrival process. This may overstate the variance of outstanding repairs somewhat, causing the recommended stock levels, which are based on the distribution of backorders, to be conservative.

We also performed an analysis of variance (ANOVA) on the mean deviation between the predicted and simulated Ao in each scenario. Effects close to zero were lumped together to create an error term. We found four effects to be significant at $\alpha = 0.05$: VMR, followed by the VMR/price interaction and the effects of price and failure rate. Experiments with a 95% target Ao revealed less information since the availabilities were close to the 100% limit. Some care should be taken with the interpretation of these results. Since the ANOVA is based on the output of a simulation model, each of the factors can eventually be made significant given enough replications. However our results do indicate the importance of the factors relative to one another.

Although we do not display the results here, in subsequent analyses we also included the availability of lateral re-supply for all three parts as a factor. The results showed that lateral re-supply improved the achieved Ao in every scenario. However, the absolute difference in Ao determined by the simulation was lower than that suggested by Vari-Metric, perhaps because the Ao predicted by Vari-Metric without lateral re-supply is conservative.

Table 2: System Ao for Different Levels of Control Variables

	High VMR (+)				Nominal VMR (–)			
	Nominal Price (+)		Low Price (–)		Nominal Price (+)		Low Price (–)	
	V-Metric	Sim	V-Metric	Sim	V-Metric	Sim	V-Metric	Sim
High Failure Rate (+)	+ + + 85.21%	+ + + 92.49% ± 0.10	+ – + 85.06%	+ – + 93.18% ± 0.08	– + + 85.13%	– + + 90.99% ± 0.11	– – + 85.64%	– – + 91.52% ± 0.07
Nominal Failure Rate (–)	+ + – 85.32%	+ + – 92.50% ± 0.10	+ – – 85.35%	+ – – 93.31% ± 0.07	– + – 85.29%	– + – 90.97% ± 0.12	– – – 86.06%	– – – 91.60% ± 0.07

Table 3: Comparison of Policies Generated by Vari-Metric and Optquest®

		Base 1	Base 2	Base 3	Base 4	Base 5	Depot
Part 1	V-Metric	4	4	9	3	6	73
	Optquest®	4	4	9	4	7	69
Part 2	V-Metric	1	1	2	1	2	7
	Optquest®	1	1	2	1	2	7
Part 3	V-Metric	1	1	3	1	2	32
	Optquest®	1	1	3	1	2	28

3.2 The Efficient Frontier

The previous analysis shows that Vari-Metric recommended stock levels may overshoot the target Ao. However we did find evidence to suggest that the Vari-Metric solutions are on the efficient frontier for the multi-criteria optimization problem of minimizing cost and maximizing Ao. Under the conditions of the base case with 85% target Ao, we used Arena's Optquest® optimization tool to search for a cheaper policy with the same estimated performance as the Vari-Metric policy. Optquest® employs a Taboo and scatter search method. Using the Vari-Metric stock levels as a starting point for the search, Optquest® analyzed 106 alternative solutions, but it did not find a less expensive solution with the same Ao.

We then used Optquest® to search for the least expensive solution that achieved an Ao greater than the original target of 85%. After considering 45 alternatives, Optquest found a stock arrangement that delivered an 86.65% Ao at a capital spares investment price of $13,573,000. This represented a $795,000 or 5.5% savings over the $14,368,000 cost proposed by Vari-Metric. The conclusion is that one may be able to achieve the target Ao with a lower inventory than that proposed by Vari-Metric, but such a policy would sacrifice a cushion over the target Ao. Table 3 displays the solution proposed by Optquest®, which is very similar to the Vari-Metric policy, effectively lowering cost and sacrificing some Ao by trimming the depot stock levels for parts 1 and 3. This solution reduced the part 1 depot level by four, but added one part to the base 4 and 5 stock levels, for a net decrease of two for part 1.

3.3 Further Analysis of VMR and Part Interactions

The analysis described in Section 3.1 indicated that VMR was the most significant of the three factors, so we performed further experiments in which we modified the 85% target Ao base case, varying the levels of VMR for part 1. Beginning with a VMR of 1.45 (the nominal value), we increased the VMR to twenty, as shown in Table 4. (We also show the results of generating part 1 failures using a standard Poisson process, which has a VMR of one.) As before we used Vari-Metric to obtain stock policies then simulated these policies to estimate actual system Ao.

The estimated Ao grew with increasing VMR, indicating that Vari-Metric was indeed overcompensating for the effects of the higher variability of part 1. Surprisingly, the growth in the overall Ao could be attributed more to growth in the availability of part 3 than part 1, which was relatively stable. We should note here that part 3 is the most expensive part, and it had the most opportunity to increase its availability. It is clear from Table 4 that while part 1's availability grew slightly with higher levels of VMR, part 3's availability grew at a much faster pace and had a larger impact on the overall Ao.

Table 5 shows the growth in the target stock levels of parts 1 and 3 (summed across locations) recommended by Vari-Metric for increasing levels of part 1's VMR. The overall target stock level for part 1 grew by 78% as part 1 VMR increased from 1.45 to 20, while the targets for parts 2 and 3 grew by 29% and 23% respectively. The rapid growth of the stock level for part 1 maintained a relatively stable availability for part 1 while the slower growth in of the stock level for part 3 caused a significant growth in part

Table 4: System Ao for Vari-Metric Policies under Different Part 1 VMRs

	Part 1 VMR										
	1 Poisson	1.45	2	3	4	5	6	8	10	15	20
Ao	90.22% ± 0.11	90.97% ± 0.12	92.06% ± 0.10	92.50% ± 0.10	92.81% ± 0.10	93.52% ± 0.09	94.12% ± 0.08	94.41% ± 0.08	94.84% ± 0.07	96.12% ± 0.07	96.69% ± 0.07
Part 1 Avail	97.70% ± 0.03	97.51% ± 0.03	97.81% ± 0.04	97.45% ± 0.05	97.79% ± 0.04	97.90% ± 0.05	97.94% ± 0.05	98.19% ± 0.04	98.20% ± 0.04	98.22% ± 0.05	98.44% ± 0.05
Part 2 Avail	99.95% ± 0.00	99.95% ± 0.00	99.96% ± 0.00	99.96% ± 0.00	99.95% ± 0.00	99.95% ± 0.00	99.95% ± 0.00	99.97% ± 0.00	99.97% ± 0.00	99.98% ± 0.00	99.99% ± 0.00
Part 3 Avail	92.57% ± 0.11	93.51% ± 0.12	94.29% ± 0.10	95.09% ± 0.10	95.07% ± 0.09	95.67% ± 0.09	96.24% ± 0.07	96.26% ± 0.07	96.68% ± 0.06	97.92% ± 0.05	98.25% ± 0.04

Table 5: Target Stock Levels for Vari-Metric Policies under Different Part 1 VMRs

	Part 1 VMR										
	1 Poisson	1.45	2	3	4	5	6	8	10	15	20
Part 1 Target	97	99	103	107	113	117	112	122	140	157	176
Part 2 Target	14	14	14	14	14	14	14	15	15	16	18
Part 3 Target	39	40	41	42	42	43	44	44	45	48	49

3's availability as well as the overall Ao. The resulting Ao was substantially higher than the 85% target.

Evidently the stock levels recommended by Vari-Metric for high VMR parts are appropriate, but the procedure seems to buy more of other parts to compensate for perceived lower availabilities in the high VMR parts. The overall effect is a net growth in system availability, above and beyond the target. For this reason the Coast Guard analysts were inclined to enforce a cap on VMR, but as the Ao grows steadily with increasing VMR, it is difficult to choose an appropriate cap. Since Vari-Metric appears to handle the stock levels for high VMR parts rather well, we suggest the following simple heuristic. Perform two runs of Vari-Metric, one with the high VMR parts set with lower variances, and one with their actual variances. Use the stock levels from the first run for the low VMR parts, and the stock levels from the second run for those with high VMR. The intuition for this heuristic is to allow Vari-Metric to set higher stock levels for the high VMR parts while preventing the upward growth in system Ao observed in Table 4 by holding the other parts to lower stock levels.

Table 6 displays the performance of various policies given an 85% target Ao and high VMR for part 1. (This was the ++- case from Table 1.) The table includes the Vari-Metric policy, a policy generated by Optquest® using the Vari-Metric policy as a starting point (minimizing cost subject to maintaining an Ao of greater than the Vari-Metric baseline), and a policy generated by our simple heuristic. For comparison we also show the performance of the Vari-Metric policy for the 85% target Ao base case (-+-).

In moving from the -+- to the ++- case, Vari-Metric raised the overall Ao roughly 1.5 percentage points by purchasing more of the higher VMR part 1 to hold that part's availability nearly stable, and by investing more in part 3 to increase that part's availability. Optquest® achieved a solution nearly as good as the -+- case in terms of Ao while saving $558,000 in parts expenditures. To do so, Optquest® invested more in part 1 to account for the higher VMR, but not as much as the ++- Vari-Metric solution. Optquest® also held the part 2 and 3 levels the same as the original -+- Vari-Metric solution. For the heuristic solution, we used the part 1 stock level from the ++- Vari-Metric policy and the parts 2 and 3 stock levels from the -+- Vari-Metric policy. The heuristic solution performed reasonably well, yielding only a slight increase in Ao over the -+- case at a lower cost than the ++- Vari-Metric solution.

3.4 The Impact of Negative Binomial Failure Counts

We also tested the effect of generating failures according to a negative binomial distribution as opposed to a Poisson process with the same mean number of arrivals. Since we were interested in the interaction of this assumption with other assumptions about the stochastic nature of the inventory system, we performed an analysis in which we also varied two other control variables. First, the Vari-Metric

Table 6: Comparison of Various Policies

	– + – Baseline V-Metric	+ + – V-Metric	+ + – Optquest	+ + – Heuristic
Parts cost	\$14.368 million	\$15.328 million	\$14.770 million	\$15.011 million
Overall Ao	90.97% ± 0.12	92.50% ± 0.10	90.49% ± 0.11	91.16% ± 0.10
Part 1 Avail.	97.51% ± 0.03	97.45% ± 0.05	97.00% ± 0.05	97.61% ± 0.10
Part 2 Avail.	99.95% ± 0.00	99.96% ± 0.00	99.96% ± 0.00	99.95% ± 0.00
Part 3 Avail.	93.51% ± 0.12	95.09% ± 0.10	93.53% ± 0.10	93.59% ± 0.10

procedure assumes that repair times are constant; in our system we estimated the repair times for parts 1, 2, and 3 to be 108, 172, and 215 days, respectively. In order to get some understanding of the consequences of this assumption, we allowed the repair times to be distributed uniformly plus or minus 30 days from the deterministic times. The Vari-Metric procedure also assumes infinite repair capacity, and we expected that the impact of deterministic repair times might be influenced by queueing effects. We therefore constructed a version of the simulation model in which parts compete for repair resources. We determined the average number of parts undergoing repair in the original model and set the number of repair resources equal to this number inflated by 10%.

The results of these experiments are displayed in Table 7. An ANOVA based on this output showed that only the assumptions concerning the distribution of the failure process and the repair capacity were significant at $\alpha = 0.05$. Clearly the extra variability associated with the negative binomial distribution reduced the operational availability from that obtained assuming a Poisson arrival process. Although a more complete analysis should be conducted, the fact that neither the distribution of repair time nor its interaction with repair capacity was significant confirms Sherbrooke's (1992) appeal to Palm's Theorem as justification for using only the mean of the repair time in Vari-Metric regardless of their distribution.

4 CONCLUSION

This study illustrates the benefits of using simulation to validate another optimization method beyond the simple matching of model output. By drilling down to individual part availability levels, the simulation revealed interesting characteristics of the Vari-Metric procedure. Knowledge of these characteristics led to a possible method for overcoming one of the shortcomings of the Vari-Metric method. By optimizing the simulation model we also gained insight concerning the efficient frontier of this multi-criteria optimization problem.

Table 7: The Impact of Assumptions Concerning the Failure and Repair Processes

	Fixed Repair Times (+)		Stochastic Repair Times (–)	
	Infinite Repair Capacity (+)	Repair Queues (–)	Infinite Repair Capacity (+)	Repair Queues (–)
Negative Binomial Failures (+)	(+ + +) 90.97% ± 0.12	(+ – +) 90.23% ± 0.11	(– + +) 91.10% ± 0.11	(– – +) 90.30% ± 0.13
Poisson Failures (–)	(+ + –) 93.70% ± 0.07	(+ – –) 93.01% ± 0.07	(– + –) 93.63% ± 0.07	(– – –) 92.94% ± 0.08

We believe the method for generating aircraft part failures from a negative binomial distribution holds promise for modeling any type of arrival process in which the distribution of the number of arrivals per time period is unimodal and the variance exceeds the mean. It is apparent that the ability to vary the VMR of failure counts had a significant impact on the simulation output of our model. This method is useful when arrival data exists in the form of number of arrivals per time period, which is common in a number of other applications.

ACKNOWLEDGMENTS

This research was made possible by funding from the U.S. Coast Guard. We especially wish to thank Mike Shirk and Jorge Salazar at the Coast Guard Aviation Repair and Supply Center in Elizabeth City, NC.

REFERENCES

Bratley, P., Fox, B.L., and L. E. Schrage. 1987. *A Guide to Simulation,* 2nd ed. New York: Springer-Verlag.

Graves, S. 1985. A Multi-Echelon Inventory Model For a Repairable Item with One-For-One Replenishment. *Management Science* 31, 1247-1256.

Law, A. and D. Kelton. 2000. *Simulation Modeling and Analysis,* 3rd ed. Boston: McGraw-Hill.

Pichitlamken, J., Deslauriers, A., L'Ecuyer, P., and A. Avramidis. 2003. Modeling and Simulation of a Telephone Call Center. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P.J. Sanchez, D. Ferrin, and D.J. Morrice, 1805-1812.

Sherbrooke, C.C. 1992. *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques.* New York: John Wiley & Sons.

Sokhan-Sanj, S., Gaxiola, G., Mackulak, G., and F. Malmgren. 1999. A Comparison of the Exponential and the Hyperexponential Distributions for Modeling Move Requests in a Semiconductor Fab. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P.A. Farrington, H.B. Nembhard, D.T. Sturrock and G.W. Evans, 774-778. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via <`www.informs-cs.org/wsc03papers/231.pdf`> [accessed July 1, 2005].

Zamperini, M. 2005. A Simulation Analysis of the Vari-Metric Optimization Method for the U.S. Coast Guard with a Part Failure Counting Process Described by the Negative Binomial Distribution. Masters thesis, Department of Industrial Engineering, Penn State University, University Park, Pennsylvania.

AUTHOR BIOGRAPHIES

MICHAEL B. ZAMPERINI is an instructor in the Department of Mathematics at the U.S. Coast Guard Academy. His research interests include simulation modeling and analysis and optimization of workforce cross-training. His e-mail address is `mzamperini@exmail.uscga.edu`.

MICHAEL FREIMER is an Assistant Professor of Operations Management in the Department of Supply Chain and Information Systems in the Smeal College of Business at Penn State University. His research interests include simulation methodology and stochastic optimization. His e-mail address is `mbf10@psu.edu`.