

FEASIBILITY STUDY OF VARIANCE REDUCTION IN THE LOGISTICS COMPOSITE MODEL

George P. Cole, III
Alan W. Johnson
J. O. Miller

Graduate School of Engineering and Management
Air Force Institute of Technology
2950 Hobson Way, Bldg 641
Wright Patterson AFB, OH 45433-7765, U.S.A.

ABSTRACT

The Logistics Composite Model (LCOM) is a stochastic, discrete-event simulation that relies on probabilities and random number generators to model scenarios in a maintenance unit and estimate optimal manpower levels through an iterative process. Models such as LCOM involving pseudo-random numbers inevitably have a variance associated with the output of the model for each run. Reducing this output variance can be costly in the additional time needed for multiple replications. This research explores the application of three different methods for reducing the variance of the model's output. The methods include Common Random Numbers, Control Variates, and Antithetic Variates. The result is a successful variance reduction in the primary output statistics of interest using the application of the Control Variates technique, as well as a methodology for the implementation of Control Variates in LCOM.

1 INTRODUCTION

The Logistics Composite Model (LCOM) is one of the Air Force's primary tools for determining optimal logistics and maintenance manpower levels. The LCOM is a stochastic, discrete-event simulation that relies on probabilities and random number generators to model scenarios in a maintenance unit by manipulating certain variables. Manpower levels are attained through an iterative process in which the variables consisting of supply, facilities, and equipment are set based on command standards. Manpower levels are adjusted after each run until a desired Sortie Generation Rate (SGR) is attained (Boyle 1990).

In a model such as LCOM, many real-life characteristics exhibit random behavior. As Law and Kelton (2000) state, "A simulation of any system or process in which there are inherently random components requires a method of generating or obtaining numbers that are *random*, in

some sense." The random number generators aid the customer in simulating the randomness of the system by producing a stream of continuous, uniformly distributed numbers between 0 and 1. The intent of the random number generator is to produce these numbers independently. However, the computer is actually using a recursive algorithm that produces numbers that seem independent, but instead follow a pattern that can be repeated over and over, called a stream (Kelton et al. 2004). These types of random number generators are called pseudo-random generators.

Models involving pseudo-random numbers inevitably have a variance associated with the output of the model for each run. The reduction of the variance in the results of the model can be costly in the form of time for multiple replications or producing a range of estimates that is too wide to realistically analyze. Simply increasing the number of replications is not always realistic for reducing variance. Replicating prior work by Bednar on the Thunder simulation, we examined the application of other methods for reducing the variance of LCOM's output, including *common random numbers*, *control variates*, and *antithetic variates* (Bednar 1996). When applied to a model such as LCOM, these variance reduction techniques may significantly reduce the variance without increasing the number of replications (Law and Kelton 2000).

2 LITERATURE REVIEW

2.1 Logistics Composite Model

The U.S. Air Force's Logistics Composite Model has existed since the late 1960s, created through a combined effort by the Rand Corporation and the Air Force Logistics Command to "relate base-level logistics resources with each other and with sortie generating capability" (Boyle 1990). While the model is capable of studying the interac-

tions between several variables, it has evolved to be known as one of the Air Force’s primary tools for establishing manpower levels in operational maintenance units and exists as part of the Air Force’s Standard Analysis Toolkit (AFSAT) (Juarez 2006).

Two separate versions of the LCOM exist in the Air Force today, one at the Air Force Manpower Agency (AFMA), Randolph AFB, Texas, and the other at the Aeronautical Systems Center’s (ASC) Systems Supportability Analysis Branch, Wright-Patterson AFB, Ohio (Dawson 2006). These two separate models essentially perform the same function, with some minor differences in the user interface (Dierker 2006). The AFMA version has four primary users: Air Combat Command, Air Mobility Command, Air Force Special Operations Command, and Air Education and Training Command. The ASC is the only primary user for the ASC version (Dawson 2006). The AFMA version is used by the MAJCOMs to derive 65-70% of their maintenance manpower requirements (Dawson 2006). The rest comes from Air Force Instructions and other guidance. The ASC version is used “to analyze manpower requirements for acquired weapon systems (as well as evaluate manpower requirement changes resulting from modifications to current weapon systems)” (Dawson 2006).

The LCOM model consists of multiple submodels, including an input model, a main model, and several post processors. The input model analyzes input data from the user and makes assumptions and corrections when necessary so that the data can be used by the main model and post processors. This data typically includes maintenance data from the Air Force’s Maintenance Data Collection systems, essential tasks needed to be performed to service each aircraft, mission requirements and flying times. The main model is the heart of the simulation, and the primary source of data for our research. It uses maintenance data and sortie data together with the process logic shown in Figure 1.

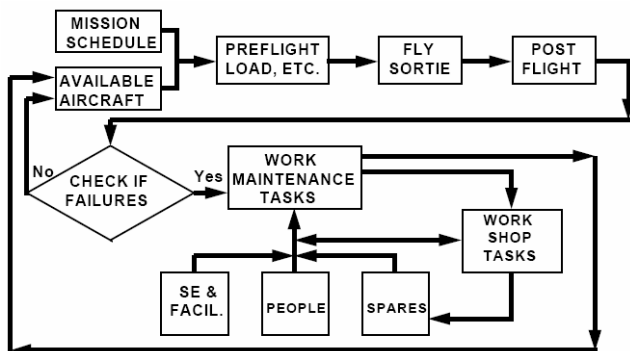


Figure 1: LCOM simulation logic (ASC/ENM 2004)

Various post processors show simulation results as a function of time, such as manpower demands, resource and

facility usage, parts availability, and depot workload (ASC/ENM 2004).

2.2 Variance Reduction Techniques

We examined three approaches for reducing LCOM’s output variance. The common random numbers approach involves comparing multiple scenarios in the same overall model. Using this approach, the individual sources of randomness are synchronized using the same random number streams across the two scenarios. Then, configurations in two different scenarios will use the same random numbers so that the different scenarios can experience similar experimental conditions (Law and Kelton 2000).

The second technique, control variates, involves identifying potential controls within the model that can be used to reduce the variance in the output. This method requires the identification of a particular random variable or variables with known expected values that are thought to correlate to an output variable, either positively or negatively. Then, using these potential controls and the estimated correlation, the expected value of the output variable is adjusted up or down based on the differences between the observed values of the control variates and their known expected values (Kelton et al. 2004).

The last technique, antithetic variates, “attempts to induce negative correlation between the results of one replication and another, and uses this correlation to reduce variance” (Kelton et al. 2004). This involves a second replication that replaces a stream’s i^{th} random number U_i with the random number $1-U_i$. For example, where U_i is used for a particular purpose, $1-U_i$ is used in the second replication for the same purpose. The pairs are averaged, possibly replicating this for several pairs (Law and Kelton 2000).

3 METHODOLOGY

3.1 Model and Output Variable Selection

Development of an actual, real-world scenario for LCOM can take many months to develop (Erdman 2006). Fortunately, the LCOM program includes two example scenarios for training and education. The first scenario, the Bicycle Model, is a very simplistic model involving a bicycle used for delivering papers early every morning. The second scenario, the Joint Service FX99 Generic Fighter Model, is loosely based on an F-16 aircraft maintenance unit, with information compiled from data at Hill AFB from July 1979 to June 1980 (ASC/ENM 2004). The generic fighter model is realistic and similar to current LCOM models in use today. Additionally, the variance in the output is sufficiently large enough for a reduction in the variance to be possible and desirable. For these reasons,

the Generic Fighter Model was selected for analysis in this study.

The original configuration of the generic fighter model is sufficient for both AV and CV, so the default configuration was used for these techniques. However, since CRN requires the comparison between two different configurations, a modification of the model for this technique was necessary. To modify the model in a way that would have a significant effect on the outcome of the output variables of interest, the manpower availability was adjusted and the model run with a 30-replication production run until a significant difference in the two output statistics of interest were empirically observed. Then, a paired difference test of hypothesis was conducted for $\mu_d = (\mu_1 - \mu_2)$ in order to verify that the mean for the two scenarios are not equal.

The period of interest for this experiment consisted of a 5-day period following a 20-day warm-up period. After consulting with LCOM users it was determined that a 20-day warm-up was likely to be sufficient for the model to exhibit a steady-state behavior, based upon their experiences and recommendations (Erdman 2006).

Although LCOM generates several dozens of various statistics, the decision was made to analyze the variance of two output statistics in particular, C15 – *Overall Achieved Sorties per Aircraft per Day*, and C24 – *Mission Capable Rate*. This decision was reached for several reasons. Since LCOM users focus almost exclusively on the C15 statistic, the reduction in the variance of any other statistic separate from C15 would have little or no impact on the way the model was run and analyzed. Additionally, *Mission Capable Rate* has a direct impact on the sortie generation rate, so C24 was also included. Second, the Generic Fighter Model produces a fair amount of variance in both output statistics C15 and C24, making variance reduction a feasible and desirable goal for our research and analysis.

3.2 Common Random Numbers

The random number generator in LCOM's programming language, SIMSCRIPT II.5, allows the user to identify up to 10 unique random number streams. In the LCOM's source code for the main model, 33 points of consumption were identified where random numbers are drawn and consumed. In the original configuration, these 33 points of consumption are divided to sample from 9 of the 10 available unique random number streams.

Each of the 33 points of consumption in the model was investigated one-by-one to determine the relative use of each point in the model where random numbers were generated and consumed. After investigating the 33 random variate draws in the LCOM main model source code, it was determined that 8 points of consumption were in use for the FX99 model. In other words, only 8 of the 33 existing points of consumption in the default settings of FX99 actively generated and consumed random numbers.

Consequently, the code was modified so that each active point of consumption was synchronized by identifying a unique random number stream to each of the points in the model. Once the source code was modified to reflect these changes, the code was compiled using Simscript's SimStudio compiling environment and an LCOM production run of 30 replications was conducted with the new, modified source code using the default scenario with unconstrained manpower and a specific seed set for the 30 replications. Then, the scenario was changed to the second scenario with constrained manpower resources and the 30-replication production run was completed again, using the exact same seed set. The paired difference from each replication was used to create the confidence interval and half-width.

Similarly, the unmodified model was run and a half-width calculated for both the unconstrained manpower scenario and the constrained manpower scenario in order to provide a base from which to compare the new results. The reduction in the confidence interval halfwidth was calculated to determine the degree of variance reduction.

3.3 Antithetic Variates

Whereas CRN requires a comparison between two different scenarios, antithetic variates relies on an induced correlation among replications within the same scenario. Since no additional scenarios are required for this technique, the default settings for the FX99 Model were used.

Fortunately, SIMSCRIPT II.5 allows the incorporation of antithetic variates with some very simple modifications. According to the manual, *Building Simulation Models with SIMSCRIPT II.5*, "To use an antithetic variate in any random deviate generator in SIMSCRIPT II.5, it is merely necessary to negate the random number stream parameter to the function `random.f`" (Russell 2000). For example, `RANDOM.F(4)` is a pseudo-random number uniformly distributed between 0 and 1, drawn from SIMSCRIPT II.5's random number stream 4. In order to implement antithetic variates, the portion of the code `RANDOM.F(4)` simply becomes `RANDOM.F(-4)` in order to turn a stream's k^{th} random number U_k into $1 - U_k$ (Russell 2000).

Like CRN, the new code is modified to reflect the change at each of the 33 points of consumption in the model, compiled in SimStudio, and run using both the unmodified code as well as the modified code with antithetic variates. Like CRN, a one-for-one relationship between U_k and $1 - U_k$ is essential to synchronize the model and induce the correlation required to achieve a reduction in variance. For this reason, careful attention must be paid to making sure the starting seeds for each stream remained the same for both the original model and the model with the antithetic variates. A production run of 15 replications was made with the unmodified code and 15 with the antithetic variates code. The output for both productions runs

were combined and averaged in order to obtain the new confidence interval calculation and halfwidth. The results were then compared to a 30-replication production run with the unmodified model.

3.4 Control Variates

The Control Variate method relies on the relationship between a random variate input and the output variable of interest. This correlation along with the deviation from the known expected value of the input variable are used to adjust the output variable up or down, closer to the true but unknown mean. In LCOM, the number of potential controls can reach the hundreds. However, each potential control must be considered by capturing each random variate value drawn in the model, along with the expected value, distribution type and parameters over the course of the 5-day period of interest.

The first step in this process involved modifying the LCOM main model source code in order to capture each random variate drawn and consumed in the model. Capturing the random variates required the use of the “print” command to dump the random variates to an output file called a PSR report for each replication in the LCOM production run. Once in the PSR report, the random variates were cut and pasted to an Excel spreadsheet where the numbers were sorted and separated by distribution type and parameter. Once sorted, each unique distribution formed the basis for a potential control to be analyzed. This means that different distribution types and parameters could come from a single random variate draw. The total number of potential controls was 23. After eliminating those potential controls with no random variates drawn in any single replication, the number of potential controls decreased to 20. Since regression theory requires at least as many replications as potential controls, 21 replications were performed (McClave et al. 2005).

The control variate calculations for each of the 21 replications were performed in Excel. Finally, the output statistics C15 and C24 for each replication were recorded. The differences between the observed and expected means for each potential control were arranged in a 20 X 21 matrix and exported to the Minitab statistical software package for the stepwise regression calculation. The predictors identified in the stepwise regression automatically became the control variates used to perform the control variate calculations, ultimately concluding with a new, smaller confidence interval and halfwidth.

4 RESULTS AND ANALYSIS

4.1 Common Random Numbers

The common random numbers model was compared to the base model using the difference between the two scenarios with constrained and unconstrained manpower as specified in section 3.1. A summary of the results for the C15 statistic, *Overall Achieved Sorties per Aircraft per Day*, with the confidence interval (CI) calculated about the mean difference from replication to replication between the two scenarios, is shown in Figure 2:

Base -- C15		
95% CI:	0.215352775	0.26398056
CI halfwidth:	0.024313892	
CRN -- C15		
95% CI:	0.166943658	0.23305634
CI halfwidth:	0.033056342	
Improvement:	-35.96%	

Figure 2: Common random numbers results summary, C15 statistic

Likewise, the same calculations were conducted for the models using the C24 statistic, *Mission Capable Rate*, as the variable of interest. A summary of the results for the C24 statistic is shown in Figure 3:

Base -- C24		
95% CI:	3.347019569	4.78031376
CI halfwidth:	0.716647098	
CRN -- C24		
95% CI:	3.825838099	5.15482857
CI halfwidth:	0.664495234	
Improvement:	7.28%	

Figure 3: Common random numbers results summary, C24 statistic

The negative improvement in Figure 2 indicates an increase in the confidence interval size, the opposite of the desired effect in this experiment. As Figure 2 shows, the confidence interval for the C15 statistic did not improve from the base model to the CRN model. In fact, the confidence interval halfwidth was significantly larger in the CRN model. On the other hand, the confidence interval halfwidth for the C24 statistic (Figure 3) improved slightly by 7.28 percent.

To put this improvement in the C24 variance into perspective, the user would need to run approximately 35 replications in order to achieve the same confidence in the C24 output statistic, versus 30 replications with the CRN model.

4.2 Antithetic Variates

The results for antithetic variates were even less favorable than the common random numbers results. In this experiment, the base, unmodified FX99 model was run a total of 30 replications. Then, the main model code was modified to incorporate the antithetic variates at each point in the model where a random variate draw occurs. The code was compiled and the FX99 model was run again for a total of 30 replications, 15 with U_j and 15 with $1-U_j$. A summary of results for the C15 output statistic is shown in Figure 4:

Base -- C15		
95% CI:	2.17456992	2.230763413
CI halfwidth:	0.02809675	
Antithetic Variates -- C15		
95% CI:	2.14564626	2.216353742
CI halfwidth:	0.03305634	
Improvement:	-25.83%	

Figure 4: Antithetic variates results summary, C15 statistic

Like the common random numbers experiment, the confidence interval for the C15 output statistic did not improve, but actually worsened by more than 25 percent. Similar results were observed in the C24 output statistic. A summary of the C24 results is shown in Figure 5:

Base -- C24		
95% CI:	66.5931921	67.21947454
CI halfwidth:	0.31314121	
Antithetic Variates -- C24		
95% CI:	66.7662962	67.59970375
CI halfwidth:	0.41670375	
Improvement:	-33.07%	

Figure 5: Antithetic variates results summary, C24 statistic

Contrary to CRN, the C.I. halfwidth for the C24 output statistic grew much larger after the implementation of the antithetic variates.

It is apparent that the common random numbers and antithetic variates synchronization techniques do not reduce the variance of the output statistics. In fact, the improvement observed in the variance of C24 in the CRN model more than likely occurs due to the randomness of the model, and not due to any true synchronization. When examining the random variate draws in the model and the purpose of each specific draw, it is easy to understand why synchronization is not possible the way the model is currently constructed. The best way to explain this is via an example. In LCOM, task durations all originally sampled from random number stream 2. We determined that the task durations were generated using three different random

variate draws in different places in the model. However, upon further examination, the vast majority of task durations are all generated and consumed from a single random variate draw. For a maintenance unit, this includes all scheduled and unscheduled task times generated for every flightline, backshop, and depot task, regardless of the priority. In the current configuration of LCOM, essentially all tasks sample from a single point in the model. If the individual tasks are not separated, the one-for-one relationship between random variates consumed across different scenarios is not achievable.

4.3 Control Variates

Unlike the previous two methods, the control variates method does not attempt to induce a correlation in the model. Instead, the control variates method attempts to take advantage of an *existing* correlation between a random variate and a particular output variable of interest.

This method was applied to two different output variables of interest, C15 – *Overall Achieved Sorties per Aircraft per Day*, and C24 – *Mission Capable Rate*. Since 20 different potential controls were initially identified, the model was run with 21 replications in order to satisfy the regression theory requirement for at least as many replications as potential predictors (McClave et al. 2005).

It is interesting to note the function of the potential controls. Controls 1-13 appear to be flightline pre- and post-sortie tasks and are all lognormally distributed. Controls 14-20 are random variates for node selection in task networks with multiple options and a specific probability associated with the task selection. Controls 14-20 are all random variates drawn using the RANDI.F function generating a uniform, discrete random integer between 1 and 100000.

After calculating the difference between the known and expected value for each potential control following each replication, the data were imported into Minitab 14 for the stepwise regression calculations, first using C15 as the response variable. An initial alpha level of 0.05 was set, and the stepwise regression returned just one control, a post-sortie task, which involves removing the missile rack from the aircraft.

After performing the control variate calculations, the new, reduced confidence interval was calculated. The improvement in C15 was dramatic, as is seen in Figure 6.

Using the control variate calculations to take advantage of the correlation between the random variate and the C15 output statistic, the confidence interval halfwidth decreased by over 56 percent. To put this improvement into perspective, the user would need to run approximately 115 replications with the original model in order to achieve the same confidence interval in the C15 output statistic, versus 21 replications with the control variates model.

Base -- C15

95% CI:	1.675734014	1.7680755
CI halfwidth:	0.046170748	
Control Variates -- C15		
95% CI:	1.708811625	1.7492334
CI halfwidth:	0.020210863	
Improvement:	56.23%	

Figure 6: Control variates results summary, C15, single control

For the second experiment with the response variable replaced by C24, the stepwise regression again determined that a single control was a predictor of the C24 response variable. In this case the predictor was a pre-sortie task, Load Chaff Dispenser. The results from this experiment are shown in Figure 7:

Base -- C24		
95% CI:	67.5334843	69.1293728
CI halfwidth:	0.79794424	
Control Variates -- C24		
95% CI:	68.3349381	68.3421545
CI halfwidth:	0.00360816	
Improvement:	99.55%	

Figure 7: Control variates results summary, C24, single control

The improvement in the confidence interval halfwidth, using the control variate as a predictor in the FX99 model, is exceptional. The results indicate an extremely strong relationship between the load chaff dispenser control and the C24 output statistic. To achieve this same level of confidence with the original model, the user would need to make over a million additional replications with the original model. This is obviously impractical, given the time constraints involved with the LCOM model runs.

While the results for the control variates experiment show a significant improvement in the confidence interval halfwidth, the stepwise regression revealed only one control with the 0.05 alpha level. Since typical, real-life LCOM models may incorporate many more variables with a much higher complexity, the possibility of a particular model possessing multiple controls is great. In order to demonstrate the technique using multiple controls, the stepwise regression experiment was performed a second time using C15 as the response variable, this time with an alpha level of 0.15. In this case the stepwise regression concluded with 5 controls identified as predictors of the C15 response variable.

The results for this experiment were not as favorable as the previous experiment with just one control, but still showed a significant improvement over the original confidence interval halfwidth. This is expected, since 85% confidence intervals are narrower even before variance re-

duction methods are applied. The results for the multiple control variate experiment are summarized in Figure 8.

Base -- C15		
95% CI:	1.674593824	1.7692157
CI halfwidth:	0.047310938	
Control Variates -- C15		
95% CI:	1.691849065	1.761883229
CI halfwidth:	0.035017082	
Improvement:	25.99%	

Figure 8: Control variates results summary, C15, multiple controls

In this case, the user would be required to make approximately 63 replications in order to achieve the same level of confidence with the original model.

5 CONCLUSION AND RECOMMENDATIONS

The common random numbers experiment exhibited no improvement when using the C15 statistic as the output variable of interest, but exhibited limited improvement when using the C24 statistic. After investigating the configuration of random variate draws in the original model, the variance of the C24 output statistic was not uniformly reduced.

Like common random numbers, antithetic variates (AV) relies on random variate synchronization in order to induce a correlation. In CRN, this correlation occurs across multiple scenarios. In AV, the correlation is induced within replications of a single scenario. The AV method is unsuccessful in LCOM for the same reason the CRN method failed – the inability to synchronize random variates in LCOM used for the same purpose. Since several random variate draws generate input for multiple purposes within the model, the possibility for synchronization within the current framework does not exist in LCOM’s current form. Consequently, the increases in the confidence intervals using the AV method do not necessarily indicate a less effective model, but rather simply an ineffective method of variance reduction.

Unlike the previous two methods, control variates capitalize on an existing correlation between random inputs and a particular output variable of interest. In this experiment using the FX99 model, control variates performed extremely well. In all cases using both the C15 output statistic and the C24 output statistic, control variates produced a significant improvement in the output variance. A significant reduction in the output variance equates to a more accurate estimate of the mean. Additionally, with a half-width goal in mind, a reduction in variance can reduce the amount of time the user spends performing additional replications in order to achieve that specified halfwidth goal. In one case, as shown in the CV experiment using the C24

output statistic, the variance reduction achieved by the control variates method was so significant that over a million replications would be necessary to achieve the same half-width with the unmodified LCOM model.

Implementing a permanent option for the control variates method in LCOM is relatively simple for a particular model. However, the controls may change from model to model. Therefore, the potential controls for each model should be captured and arranged in a consistent manner, listing potential controls and the output variable of interest. This can be performed fairly easily by a postprocessor. The postprocessor must first capture the potential random variate controls and their known expected values, and then subtract observed means from the respective expected values to obtain the differences. Then, the regression would be performed in order to identify the significant predictors of the response variable. This regression could be performed by a postprocessor or a statistical software package such as Minitab.

ACKNOWLEDGEMENTS

The authors wish to thank the Aeronautical System Center's Systems Supportability Analysis Branch staff for their sponsorship of our research. In particular, the authors wish to thank Mr. Frank Erdman and Mr. Greg Dierker for their outstanding support and expertise.

REFERENCES

- Aeronautical Systems Center (ASC/ENM). 2004. *ASC LCOM 2.6 Users Manual*. Wright-Patterson AFB, OH: ASC/ENM, February 2004.
- Bednar, E. M. 1996. *Feasibility Study of Variance Reduction in the Thunder Campaign-Level Model*. MS Thesis.
- Boyle, E. 1990. *LCOM Explained: Interim Technical Paper for Period May 1990 – June 1990*. AFHRL-TO-90-58. Brooks AFB TX: Air Force Systems Command, July 1990 (AD-A224497).
- Dawson, K. 2006. *LCOM Process Reengineering*. Air Force Logistics Management Agency, Maxwell AFB, AL, March 2006.
- Dierker, G. J. 2006. LCOM Integrated Product Team, Systems Supportability Analysis Branch, Modeling, Simulation, and Analysis Division, Engineering Directorate, Aeronautical Systems Center, Wright-Patterson AFB OH. Personal Interviews.
- Erdman, F. J. 2006. Aeronautical Systems Center LCOM Group Lead, Systems Supportability Analysis Branch, Modeling, Simulation, and Analysis, Division, Engineering Directorate, Aeronautical Systems Center, Wright-Patterson AFB, OH. Personal Interviews. April – November 2006.
- Juarez, F. 2006. LCOM Model Description. Air Force Management Agency, Randolph AFB, TX, 2006.
- Kelton, W. D., R. P. Sadowski, and D. T. Sturrock. 2004. *Simulation With Arena*. 3rd Edition. New York: Marcel Dekker, Inc.
- Law, A. M. and W. D. Kelton. *Simulation Modeling and Analysis*. 3rd Edition. New York: McGraw-Hill, 2000.
- McClave, J. T., P. G. Benson, and T. Sincich. 2005. *Statistics for Business and Economics*. 9th Edition. New Jersey: Prentice Hall.
- Russell, E. C. 2000. *Building Simulation Models With Simscript II.5*. California: CACI Products Co.

AUTHOR BIOGRAPHIES

GEORGE P. COLE, III is an aircraft maintenance officer in the U.S. Air Force. He is a graduate of the Air Force Institute of Technology's Logistics Management Program in 2007. He is currently assigned to Air Mobility Command Headquarters, Scott AFB, IL. His experience includes Strategic Airlift and C-17 flightline operations. His e-mail address is georgecole3@hotmail.com.

ALAN W. JOHNSON is the deputy department head and an associate professor in the Department of Operational Sciences, Air Force Institute of Technology (AFIT). He has a B.S. in Mechanical Engineering from Montana State University, an M.S. in Systems Management from AFIT, and a Ph.D. in Industrial and Systems Engineering from Virginia Tech. His research interests include all aspects of military logistics, but emphasize reliability and maintainability and the logistics support aspects of space flight systems. He is a retired Air Force officer. His email address is alan.johnson@afit.edu

J. O. MILLER is a 1980 graduate of the U.S. Air Force Academy (USAF) and retired from the Air Force as a Lt. colonel in January 2003. In addition to his undergraduate degree from USAFA, he received an MBA from the University of Missouri at Columbia in 1983, his M.S. in Operations Research from the Air Force Institute of Technology (AFIT) in 1987, and his Ph.D. in Industrial Engineering from The Ohio State University in 1997. Dr. Miller is Director of the Center for Operational Analysis and an associate professor in the Department of Operational Sciences at AFIT. His research interests include combat modeling, computer simulation, and ranking and selection. He can be reached at john.miller@afit.edu.