Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

A DESIGN OF EXPERIMENTS APPROACH TO READINESS RISK ANALYSIS

Keebom Kang
Kenneth H. Doerr

Graduate School of Business & Public Policy
Naval Postgraduate School
Monterey, CA 93943-5103, U.S.A.

Susan M. Sanchez

Operations Research Department and
Graduate School of Business & Public Policy
Naval Postgraduate School
Monterey, CA 93943-5219, U.S.A.

ABSTRACT

We develop a simulation model to aid in identifying and evaluating promising alternatives to achieve improvements in weapon system-level availability when services for system components are outsourced. Two outcomes are valued: improvements in average operational availability for the weapon system, and reductions in the probability that operational availability of the weapon system falls below a given planning threshold (*readiness risk*). In practice, these outcomes must be obtained through performance-based agreements with logistics providers. The size of the state space, and the non-linear and stochastic nature of the outcomes, precludes the use of optimization approaches. Instead, we use designed experiments to evaluate simulation scenarios in an intelligent way. This is an efficient approach that enables us to assess average readiness and readiness risk outcomes of the alternatives, as well as to identify the components and logistics factors with the greatest impact on operational availability.

1 INTRODUCTION

Performance Based Contracts are becoming increasingly popular in both the Department of Defense and the commercial defense industry. Performance Based Logistics (PBL) contracts are a type of performance based contract intended to improve weapon system availability at a reduced cost.

The unique aspect of performance based contracts is their focus on outcomes; the client organization specifies key performance goals, and allows the vendor to determine the best way of obtaining those goals (ASN-RDA 2003). Such contracts are called *contra proferentem,* because in contrast to typical contract law, ambiguities in the contract (in particular, lack of detail in methods for obtaining the contracted results) are construed in favor of the client organization, rather than the vendor. Indeed, the main point of performance based contracts is to outsource not only the tasks involved in obtaining an outcome (e.g., the inventory management required to improved system availability), but also the risk associated with those tasks. In other words, the client wishes to rely on the outcomes specified in the contract, and to have the vendor bear the risks associated with insuring the delivery of those outcomes. Hence, in such contracts it is important for the client organization to evaluate not only expected outcomes, but also the associated risk (Doerr, Lewis and Eaton 2005).

One valued outcome is high system operational availability, or the average percentage of assets which are available for operations (Ao). Another is low *readiness risk* (Kang et al. 2005) which is a measure of the risk associated with contract performance. This is the probability that a vendor will fail to deliver a desired threshold of operational availability, such as the probability that less than 80% of a given type of aircraft will be available for operations at any given time. Our simulation model allows us to consider both measures of performance.

The simulation approach we describe in this paper is intended to help decision makers develop the most effective alternatives for increasing the average operational availability and reducing readiness risk of a weapon system. The alternatives involve specifying component-level outcomes for one or more of four logistics elements: component-level inventory service level, reductions in component failure rate, increases in component repair rate, or reductions in component logistics delay (the time required for transportation and administrative work). Our model captures the joint effect of all of these component-level logistics elements on operational availability, and calculates a lifecycle cost for each alternative. We then use a design-of-experiments approach developed for large-scale simulation experiments (Kleijnen et al. 2005) to sample the state space of possible alternatives in an intelligent way. Using this approach, we can estimate which logistics elements and which components have the greatest potential to improve availability.

The contribution of our work lies in the integrative nature of our solution approach. We apply a recently developed method for sampling in large-scale simulation experiments, and use a performance metric (readiness risk) designed for performance-based agreements.

2 BACKGROUND

2.1 Performance Based Logistics

There is a small but growing literature on various aspects of Performance Based Logistics (PBL) contracting. Berkowitz et al. (2003) conduct a survey of military applications of PBL, and formulate a set of best practice recommendations. Apgar and Keane (2004) describe the strategic goals of PBL, and assert that the principle of specifying outcomes rather than methods is consistent with a long-standing military strategy known as "commander's intent." Doerr, Lewis and Eaton (2005) examine metrics for PBL, and developed an argument for the centrality of risk measurement in such contracts. Kim, Cohen and Netessine (2006) look at a situation in which a contractor awarded a system-level prime contract for availability improvement must negotiate with subcontractors to achieve given component-level performance. But a recent Government Accountability Office report (GAO, 2004) is critical of systems-level PBL contracts, and recommends greater emphasis on PBL contracts at the component level to better maintain control over costs and performance. As Kang et al. (2005) show, the proper valuation and management of such component level contracts entails the development of a comprehensive model which incorporates key performance dimensions of *all* critical components. They demonstrate tradeoffs between readiness risk and lifecycle cost on given alternatives, with a numerical analysis using two (disjoint) models.

Risk-based capacity models such as the one proposed in this paper have been the subject of a great deal of research in the commercial sector (Van Miegham 2003) and have also been applied to the acquisition of production capacity for airfoils used in military aircraft (Prueitt and Park 2003). Risk-based capacity models deal with technological, demand, or price uncertainty, and are not directly applicable to the valuation of logistics services or the impact those services will have on system availability. The probability that operational availability (Ao) will remain above a certain planning threshold, or target readiness, is what we call *readiness risk*. This measure is not new—it is one of many imbedded in a system used by the U. S. Air Force for planning levels of spare-parts inventory (Slay et al. 1996). Methodologically, it is simply a type of quantile analysis. But from the war fighter's point of view, this risk may be the key performance dimension (Eaton, Doerr and Lewis 2006). The war fighter, after all, is less concerned with the average number of mission capable aircraft than with the probability that he will have enough aircraft to fly a particular mission.

Performance based contracting changes the way risk should be valued and measured in component-level contracts to improve system availability. The impact of variance in component level reliability (failure rates) and maintainability (repair time) on average system availability was well understood (Blanchard, Verma and Peterson 1996) before PBL contracts ever became popular. More recent work examines alternatives for reliability or maintenance improvement at the component level, with the primary outcome being system level availability; Cassady, Pohl, and Jin (2004) use a cost function which assumes a continuous range of available alternatives for both reliability and maintenance. However, they do not examine logistics delay, which we will show to be a critical logistics element in determining system availability, nor do they use readiness risk as an outcome measure.

Within the field of reliability engineering, *reliability allocation methods* seek to minimize the cost of allocating resources for component-level reliability in order to obtain a given system-level reliability requirement (Kececioglu 1991, pp. 363-399). These procedures generally assume a continuous range of reliability is available for each component, and that the cost of achieving higher reliability levels increases exponentially. Our work differs from this in that they are primarily focused on reliability (failure rates) as an outcome measure at the component and system level.

2.2 Design of Experiments

Clearly, simulation models of even relatively simple logistics systems can have a very large number of inputs—many of which may be uncertain or unknown—that potentially impact the model's performance. In the design of experiments (DOE) literature, these are referred to as *factors*. Factors can be qualitative or quantitative. They can include distributional models (e.g., the use of exponential, triangular, or (truncated) normal distributions for service times); parameters of these distributions (e.g., means, standard deviations, or rates); or different policy choices that determine how a subsystem within the model behaves (e.g., use of priority queues to process critical components more rapidly).

In real-world experiments, it is difficult to control more than a handful of factors at a time. This is not the case for simulation experiments, where the analyst has the ability to specify the *levels* (values) for all of the input factors before running the simulation. Still, once the factors and potential levels have been determined, this creates a huge number of potential *scenarios* (or *design points*). For example, if an analyst wished to explore nine factors, each at 10 levels, there are one billion (10^9) different scenarios that could be considered. The design might need to be replicated for stochastic simulations, because specifying all input factors does not remove randomness from the output. Such a large experiment is clearly impractical. Even if it were possible to run all scenarios in a reasonable amount of time, the volumes of output data would easily overwhelm most post-processing analytic tools, leaving the analyst limited in their abilities to statistically interpret the results.

Fortunately, efficient experimental designs can be used to specify a small number of suitable scenarios. The following characteristics are desirable (Cioppa, Lucas and Sanchez 2004; Kleijnen et al. 2005):

- The ability to examine many variables (ten or more) efficiently;
- Approximate orthogonality between inputs, to facilitate response surface metamodeling;
- Minimal *a priori* assumptions about the response surface;
- Flexibility to allow for the estimation of many effects, interactions, thresholds, and other features of the response surface; and
- An easy method for generating the design.

Kleijnen et al. (2005) discuss situations where various classes of designs are appropriate, but there is no one-fits-all design. In our explorations of readiness risk we want to screen many variables for importance, while simultaneously maintaining the ability to fit complex meta-models to a handful of input variables that are found to have the most impact on the responses. Given this, and the above design goals, the nearly orthogonal Latin hypercubes constructed by Cioppa and Lucas (2006) are particularly useful.

Designed experiments for simulation models involving many factors have been successfully applied to a host of other military applications. Links to over 40 M.S. theses by students at the Naval Postgraduate School are available online at the *SEED Center for Data Farming* website at <http://diana.cs.nps.navy.mil/seedlab>, along with links to papers, software, spreadsheets, and other tools to facilitate experimental design. Summaries of successful studies conducted in the U. S. or in several allied countries are available at the Project Albert web site at <http://www.projectalbert.org>.

3 CASE STUDY

We use the decision environment of Kang et al. (2005) in this paper, but we develop an integrative model to investigate potential alternatives for development. We are interested in readiness analyses of an unmanned aerial vehicle (UAV) squadron that has 40 aerial vehicles (AV). When a critical component in an AV fails, the faulty component is removed from the AV, an RFI (ready-for-issue) spare is installed, and the faulty component is sent to the repair facility. After the repair is complete, the component becomes an RFI spare and is sent to the spare pool. When a critical component fails, and an RFI spare is not available, the AV will be grounded (and will become not mission capable) until an RFI component is available. A failure of a non-critical component may degrade readiness, but the system is assumed to be operable (that is, mission capable or partially mission capable). In this case study, we do not consider "cannibalization," the swapping of a working component from one downed AV to another.

Our simulation model estimates the average operational availability and the readiness risk at various thresholds of interest. Our goal is to better understand how changes in reliabilities, number of spare parts and other logistics factors (e.g., repair times and transportation delays) affect the average operational availability and the readiness risk of the squadron.

We consider three critical components in this case study: engines, propellers, and avionic computers. We assume that the time between failures for each component follows an exponential distribution. The ranges of MTBF (mean time between failures) of the individual components are provided in Table 1, along with the ranges of the number of spare components, component repair times (in hours), and the transportation/logistics delay (in days).

Several designs are possible, but we use an NOLH with 257 runs (Cioppa and Lucas 2006). This design is capable of handling up to 29 factors without increasing the number of scenarios. It can be easily constructed by entering the low and high values in Table 1 into a spreadsheet (Sanchez 2006). (We remark that that ten input factors could be examined using a NOLH with as few as 33 scenarios if the time required for 257 runs was prohibitively long.) Because our model runs quickly, we opt for a larger design to allow a more detailed investigation of our model's behavior. The input parameters for the first ten scenarios are shown in Table 2. In all, there are ten different simulation inputs used as factors for our designed experiment. In addition, there is a stochastic element that occurs due to the pseudo-random numbers generated for stochastic failure times, repair times, and transportation/administrative delay times.

For each scenario, the simulation model reads a row of data from the spreadsheet excerpt in Table 2. The MTBFs of three components are first read, followed by the number of spares for each component, the means of the component repair times, and the mean for the transportation/administrative delay. The repair times are assumed to follow symmetric triangular distributions with lower and upper bounds of 0.5 (mean) and 1.5 (mean), respectively. The same approach is used for the repair time distributions. The transportation and administrative delay (in days) follows a symmetric triangular with lower and upper bounds of 0.75 (mean) and 1.25 (mean), respectively. Flight operations are conducted 24 hours per day, seven days per week. Each air vehicle operates an average of four hours per day. The repair shop operates eight hours per day, seven days per week.

Table 1. Ranges of Input Parameters

Input Parameter	Range
MTBF of Engine	200 - 400 hrs
MTBF of Propeller	150 – 300 hrs
MTBF of Avionic Computer	300 – 600 hrs
Spare engines	1 – 20 units
Spare propellers	1 – 20 units
Spare avionic computers	1 – 20 units
Repair time for engines	1 – 30 hrs
Repair time for propellers	1 – 30 hrs
Repair time for avionic computers	1 – 30 hrs
Transportation/administrative delay for each failure	1 – 15 days

4 RESULTS

We run a total of 257 scenarios, each of which is simulated over a period of 1,000,000 hours—sufficiently long that we need not be concerned about initial bias. The results of the simulation are the average Ao (operational availability) and the quantiles (10%, 20%, … , 80%, and 90%) of Ao; these are automatically written onto an EXCEL spreadsheet worksheet and then imported into the JMP® (SAS 2002) for further analysis. We remark that the outputs must be matched to the scenarios (specifically, the levels of each input factor must be available) in order to analyze the data. Also, for large experiments it can be very helpful to automate the process of running the simulation for different scenarios; see (Kleijnen et al. 2005) or (Sanchez 2006) for further discussion.

For demonstration purposes, we present only the results for the average Ao and its 80% quantile (i.e., the probability that the Ao goes below 80%). Our intent is to illustrate the types of insights that can be gained from a designed experiment approach, rather than to make inferences regarding readiness risk for a real weapons system.

4.1 Average Operational Availability

We begin assessing the output by looking at histograms of the simulation responses. This can be a way of "accidentally" performing verification and validation of a simulation model. It may reveal combinations of input factor settings for which the model does not work properly; present results that may, at first glance, challenge the analyst's intuition; or suggest additional features that should be included in the simulation model (Kleijnen et al. 2005). Our results indicate that the average operational availability differs widely across the different scenarios, ranging from 0.599 to 0.976. The average Ao across the 257 scenarios is 0.795 with a standard deviation of 0.085. It appears that at least one of the input factors does, indeed, have a substantial influence on the system's performance.

After confirming that the results appear reasonable, we turn to our main goals—identifying those factors and components that have the greatest impact on performance. A useful non-parametric tool is a regression tree, as shown in Figure 1. These graphics are effective for understanding and communicating the results of thousands of runs over many factors. Regression trees are more human-readable and can be easier to describe than multiple regression models because they reveal the structure in the data in a simple way. Initially, the data are grouped in a single cluster. All potential input factors are examined to identify how best to split them to yield two leaves so that the variability in the response within each leaf decreases and the variability in the response between the leaves increases.

Table 2. Input Parameter Settings for First Ten of 257 Scenarios

Scenario	MTBF Engines	MTBF Props	MTBF AvComp	Spare Engines	Spare Props	Spare AvComps	Mean Engine Repair (hrs)	Mean Prop Repair (hrs)	Mean AvComp Repair (hrs)	Mean Trans/Admin Delay (days)
1	280	282	478	13	13	18	30	26	6	2
2	223	210	552	19	15	11	29	25	17	10
3	232	239	335	12	19	16	28	27	20	6
4	281	174	420	13	15	15	26	29	5	9
5	273	234	484	9	19	16	23	21	3	12
6	288	205	587	4	17	18	22	29	22	7
7	209	242	432	3	11	11	23	18	20	9
8	277	156	410	9	19	15	21	21	15	2
9	215	227	579	13	8	16	22	29	6	6
10	297	161	559	15	2	13	17	17	29	12

Figure 1 shows the regression tree for predicting the average Ao from the 257 simulation scenarios. The dominant factor is clearly the average transportation/administrative delay. For example, the first split at the top indicates that the average Ao is 0.737 across the 138 scenarios that have a mean transportation/ administrative delay of eight or more days. In contrast, the average Ao is 0.862 (17% higher) among the 119 scenarios that had a mean transportation/administrative delay of less than eight days. Even with only four splits, the regression tree achieves an R^2 value of 0.74. For larger trees with many leaves, it may be helpful if the leaves corresponding to favorable, intermediate, and unfavorable outcomes are colored green, yellow, and red, respectively (Cioppa, Lucas and Sanchez 2004).

Regression trees are non-parametric approaches for fitting a statistical model to the simulation output. They can be good at identifying subsets of the output that behave much differently than the rest. Regression metamodels can also be valuable. They may confirm the regression tree results concerning which factor or factors have the greatest influence on the results, or they may allow more succinct descriptions of the simulation model's performance if it can be well-described by simple polynomial metamodels.

Accordingly, we also fit regression metamodels of the Ao as a function of main effects, quadratic effects, and two-way interactions of the ten input factors. There are a total of 65 potential terms in the model (ten main effects, ten quadratic effects, and 45 two-way interactions). We use stepwise regression to identify the most important factors, then simplify the model even further by eliminating a few terms with p-values an order of magnitude higher than the others. Our final metamodel is shown in Figure 2. The adjusted R^2 is 0.97, showing that the regression metamodel does an excellent job of explaining the variability in the simulation output. We tried other models as well. For example, a simpler model with only six significant main effects (three MTBFs, the transportation/ administrative delay, and mean repair times for the two least reliable components: propellers and engines) yields an R^2 of 0.92. This might also be used to make inferences.

The large |t_ratio| for the mean transportation/administrative delay (Figure 2) shows it to be the dominant factor, and agrees with our regression tree results. Note that the numbers of spare parts do not appear in the model. This means that raising them from their lowest levels to the highest levels in Table 1 does not lead to any appreciable improvement in the average operational availability. This suggests that it might be possible to entirely eliminate the spare parts without adversely affecting operational availability. Of course, such a possibility would need to be confirmed by running new scenarios and observing the output.

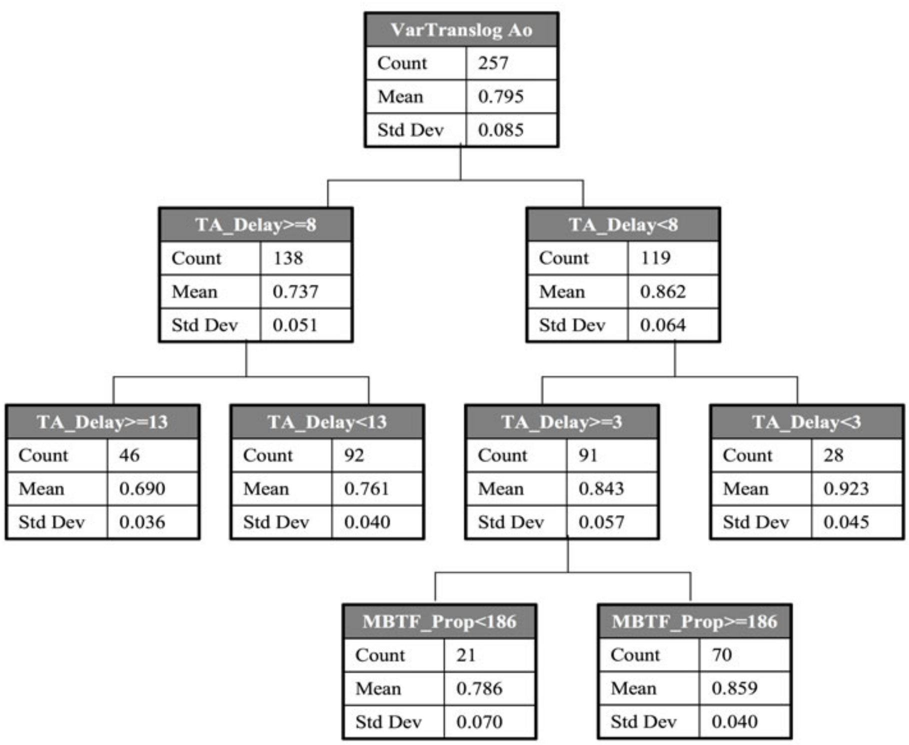

Figure 1. Regression Tree for Average Ao (First Experiment)

Response Avg Op Av				
Parameter Estimates				
Term	**Estimate**	**Std Error**	**t Ratio**	**Prob>|t|**
Intercept	0.755893	0.009407	80.36	<.0001
MTBF_ Eng	0.0002475	0.000016	15.00	<.0001
MTBF_ Prop	0.0005777	0.000022	26.26	<.0001
MTBF_ AvComp	0.0000922	0.000011	8.38	<.0001
Eng Repair hrs	-0.000762	0.000114	-6.71	<.0001
Prop Repair hrs	-0.002121	0.000114	-18.67	<.0001
Trans/Admin Delay	-0.017848	0.000234	-76.18	<.0001
(MTBF_ Eng-300.016)*(Eng Repair hrs-15.5019)	0.0000114	0.000002	5.54	<.0001
(MTBF_ Prop-225.004)*(Prop Repair hrs-15.5019)	0.0000365	0.000003	12.79	<.0001
(Prop Repair hrs-15.5019)*(Trans/Admin Delay -8.00389)	0.0002673	0.00003	8.86	<.0001
(MTBF_ Eng-300.016)*(MTBF_ Eng-300.016)	-0.000001	3.2e-7	-3.46	0.0006
(MTBF_ Prop-225.004)*(MTBF_ Prop-225.004)	-0.000004	5.9e-7	-7.42	<.0001
(Prop Repair hrs-15.5019)*(Prop Repair hrs-15.5019)	-0.000098	0.000015	-6.35	<.0001

Figure 2. Regression Metamodel for Average Ao (First Experiment)

A plot of the residuals vs. the predicted values (not shown) indicates that there are a few outliers in this metamodel. Three points result in substantially lower operational availability than predicted. Depending on the vendor PBL contract, these could be worth a closer look.

As it can be difficult to look at a regression equation and get a good sense of how the factors and interactions affect the response, *interaction plots* are often useful. The interaction plot for our regression metamodel appears in Figure 3. This interaction plot consists of several small subplots that indicate how the predicted performance (Ao) varies as a function of pairs of input factors. For example, the subplot that appears at the center of the top row shows the joint effect of the MTBF for aircraft engines and the (mean) engine repair hours. The flat upper line (in blue) shows that when the MTBF is 400 hours, changing the engine repair time between its low and high values (1-30 hours) has little impact on Ao. But if the MTBF for engines is only 200 hours (lower line, in red) then longer engine repair times decrease Ao. The difference in slopes indicates an interaction between engine MTBF and repair times: the impact of high repair times is mitigated by large MTBF. An even stronger interaction is observed between MTBF and repair times for the propellers.

Transportation/administrative delays are so dominant that we rerun the experiment after fixing the average delays to five days for all scenarios. (Note that individual delays still follow a random distribution.) These results let us focus on the other factor effects and interactions. A portion of the regression tree, corresponding to the better outcomes, is provided in Figure 4. Here, we see the impact of the MTBF and repair times for the least reliable component (propellers); the next component to show up in the tree is the engine, via its MTBF. The left-hand portion of this regression tree (not shown) has the same variables at each branch, although the "splits" at the branches occur at different factor levels.

4.2 Readiness Risk: 80th Percentile

The analyses for the 80th percentile of readiness risk are similar. The same set of factors appear as the most important determinants of performance as they did for the average operational availability, although the specific coefficients vary. When the mean transportation / administrative delay varies between one day and 15 days, it is the dominant factor in both the regression tree and the regression metamodel. The "splits" which the regression tree uses to break this delay into different components differ slightly from those for the average Ao. For example, the best leaf for readiness risk of 80% or better corresponds to an average transportation/administrative delay of less than six days and a MTBF for propellers of at least 201 hours. The best leaf in the regression tree for average operational availability corresponds to an average transportation/administrative delay of less than three days and a MTBF for propellers of at least 186 hours. These differences confirm that both measures should be considered—they are not substitutes for one another. Our regression tree with four splits and five leaves yields an R^2 of 0.78, and our regression model with seven terms (five main effects and two interactions) yields an R^2 of 0.92.

For the second experiment with the transportation/administrative delay fixed to five days, we once again find that the least reliable components are the major determinants of performance. The results are similar to those for average Ao, although the levels at which splits occur in the regression tree, and individual regression coefficients, differ.

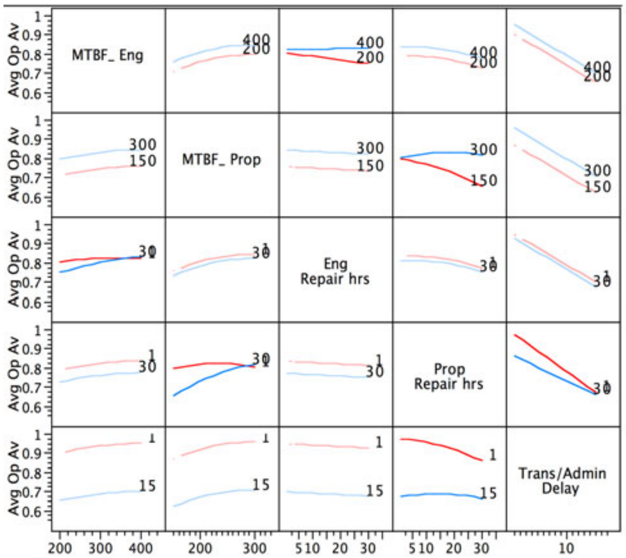

Figure 3. Interaction Profile Plot (First Experiment)

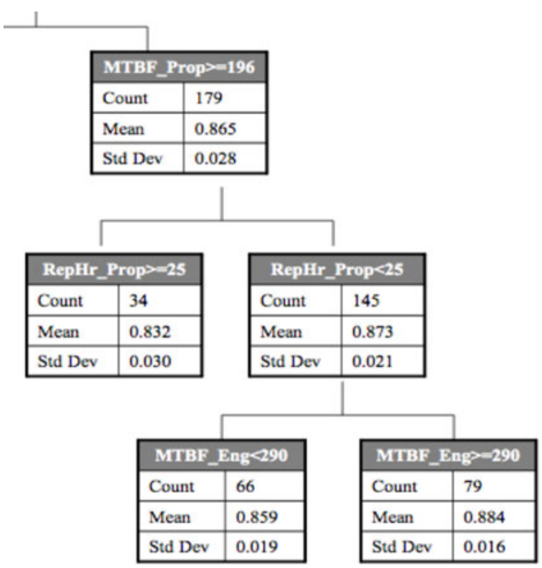

Figure 4: Partial Results for Ao (Second Experiment)

5 REMARKS

As we discussed earlier, the simulation model used in this paper is not intended to provide detailed insights regarding a particular real-world situation. For example, the use of exponential times between failures may not be appropriate, and the triangular service time distributions are unlikely to be accurate representations of real-world data. However, the same approach can easily be applied to simulations that are more realistic.

For our model, we chose to make one extremely long run for each scenario and report the aggregated results. Other alternatives are possible. The basic design (i.e., set of scenarios) could be replicated several times, or a batch means approach could be used. Shorter runs would necessitate deleting the initial transient period before computing the performance statistics, unless the performance measures of interest are the average operational availability or readiness risk for a particular fleet over a fixed period of time, rather than the steady-state results.

ACKNOWLEDGMENTS

The research was supported by NAVSEA PEO SHIPS, PEO IWS (Integrated Warfare Systems) via the Naval Postgraduate School Acquisition Research program, and the U. S. Marine Corps' Project Albert effort.

REFERENCES

Apgar IV, M. and J. M. Keane. 2004. New business with the new military." *Harvard Business Review* 82 (9): 45-56.

ASN-RDA [Assistant Secretary of the Navy – Research, Development and Acquisition]. 2003. *Performance Based Logistics Guidance Document*. Unpublished manuscript.

Berkowitz, D., J. N. D. Gupta, J. T. Simpson, J. McWilliams, L. Delayne, B. Brown, D. Cameron, and T. Sparks. 2003. *Performance Based Logistics*. Center for the Management of Science and Technology, Huntsville, AL.

Blanchard, B. S., D. Verma, and E. L. Peterson. 1995. *Maintainability: A Key to Effective Serviceability and Maintenance Management*. John Wiley & Sons, New York.

Cassady, C. R., E. A. Pohl, and S. Jin. 2004. Managing availability improvement efforts with importance measures and optimization. *IMA Journal of Management Mathematics* 15 (2): 161-174.

Cioppa, T. M. and T. W. Lucas. 2006. Efficient nearly orthogonal and space-filling Latin hypercubes. *Technometrics*, forthcoming.

Cioppa, T. M., T. W. Lucas, and S. M. Sanchez. 2004. Military applications of agent-based simulations. *Proceedings of the 2004 Winter Simulation Conference*, eds. R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 171-180. Piscataway, NJ: Institute of Electrical and Electronic Engineers.

Doerr, K. H., I. A. Lewis, and D. R. Eaton. 2005. Measurement issues in performance based logistics. *Journal of Public Procurement* 5 (2): 164-186.

Eaton, D. R., K. H. Doerr, and I.A. Lewis (in press). Performance based logistics: A warfighting focus," *U. S. Naval Institute Proceedings*.

Ebeling, C. E. 1996. Reliability and maintainability model (RAM). Report NASA-CR-203254, National Aeronautics and Space Administration Langley Research Center.

GAO [Government Accountability Office]. 2004. Opportunities to enhance the implementation of performance based logistics. Report of the Subcommittee on Readiness and Management Support, Committee on Armed Services, U. S. Senate. August 2004.

Kang, K., K. H. Doerr, M. Boudreau, and U. Apte. 2005. A decision support model for valuing proposed improvements in component reliability. Sponsored report NPS-LM 05-008, Naval Postgraduate School, Monterey, CA.

Kececioglu, D. 1991. *Reliability Engineering Handbook, V. 2.* Prentice-Hall, Englewood Cliffs, NJ.

Kim, S. H., M. A. Cohen, and S. Netessine. 2006. Performance contracting in after-sales service supply chains. Working Paper, Wharton School of Business, University of Pennsylvania.

Kleijnen, J. P. C., S. M. Sanchez, T. W. Lucas, and T. M. Cioppa. 2005. A user's guide to the brave new world of designing simulation experiments. *INFORMS Journal on Computing* 17 (3): 263-289 (with online companion).

Prueitt, G. C. and C. S. Park. 1997. Phased capacity expansion – using continuous distributions to model prior belief. *The Engineering Economist* 42 (2): 91-110.

Sanchez, S. M. 2006. Work smarter, not harder: Guidelines for designing simulation experiments. *Proceedings of the 2006 Winter Simulation Conference,* eds. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, forthcoming, Piscataway, NJ: Institute of Electrical and Electronic Engineers.

Sanchez, S. M. and T. W. Lucas. 2002. Exploring the world of agent-based simulations: Simple models, complex analyses. *Proceedings of the 2002 Winter Simulation Conference*, eds. J. L. Snowdon, J. Charnes, C.-H. Chen, and E. Yücesan, 116-126. Piscataway, NJ: Institute of Electrical and Electronic Engineers.

SAS. 2002. JMP User's Guide, Version 5: SAS Institute Inc. Cary, NC.

Slay, F. M., T. C. Bachman, R. C. Kline, T. J. O'Malley, F. L. Eichorn, and R. M. King. 1996. Optimizing spares support: The aircraft sustainability model. Report AF501MR1, LMI Corporation, McClean, VA.

Van Miegham, J. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing and Service Operations Management*, 5 (4): 269-302.

AUTHOR BIOGRAPHIES

KEEBOM KANG is the Academic Associate (Chair) of the Material Logistics and Supply Chain Management Program at the Graduate School of Business and Public Policy, Naval Postgraduate School. He received his Ph.D. from Purdue University. His research interests are in the areas of simulation modeling, and logistics applications. He was the Co-Editor of the *Proceedings of the 1995 Winter Simulation Conference*, and the program chair for the 2000 Winter Simulation conference. He is currently the Associate Editor for *Naval Research Logistics*. His e-mail is <kkang@nps.edu>.

KENNETH H. DOERR, B.S. Indiana University, 1984; Ph.D. University of Washington, 1994. Prior to joining the faculty at the Naval Postgraduate School, he taught at the University of Miami, the University of Washington and Santa Clara University, and also held research fellowships at the University of Waterloo and the University of Cincinnati. He has worked for Shell Oil, Monsanto Corporation and Peoplesoft in manufacturing and supply chain systems. His research has appeared in several leading journals, including *Management Science*, *The Academy of Management Review, IIE Transactions* and *The Journal of Applied Psychology*. His research interests are in process design, work design, and information technology to facilitate operations. Dr. Doerr is originally from St. Louis, Missouri. His e-mail is <khdoerr@nps.edu>.

SUSAN M. SANCHEZ is a Professor of Operations Research at the Naval Postgraduate School, where she holds a joint appointment in the Graduate School of Business & Public Policy and serves as the Director of the *SEED Center for Data Farming*. Her research interests include experimental design for simulation studies, data-intensive statistics, and robust selection. She has a Ph.D. in Operations Research from Cornell University. She is currently the Simulation Area Editor for the *INFORMS Journal on Computing* and the ASA representative to the WSC Board of Directors. Her e-mail is <ssanchez@nps.edu> and her web address is <http://www.nps.navy.mil/orfacpag/resumePages/sanchs.htm>.