University of Nebraska - Lincoln DigitalCommons@University of Nebraska - Lincoln

U.S. Air Force Research

U.S. Department of Defense

2002

Computing small-fleet aircraft availabilities including redundancy and spares

Jeffery K. Cochran Arizona State University

Theodore P. Lewis United States Air Force

Follow this and additional works at: https://digitalcommons.unl.edu/usafresearch

Part of the Aerospace Engineering Commons

Cochran, Jeffery K. and Lewis, Theodore P., "Computing small-fleet aircraft availabilities including redundancy and spares" (2002). *U.S. Air Force Research*. 19. https://digitalcommons.unl.edu/usafresearch/19

This Article is brought to you for free and open access by the U.S. Department of Defense at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in U.S. Air Force Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.



Computers & Operations Research 29 (2002) 529-540



www.elsevier.com/locate/dsw

Computing small-fleet aircraft availabilities including redundancy and spares

Jeffery K. Cochran^{a,*}, Theodore P. Lewis^b

^aDepartment of Industrial and Management Systems Engineering, Arizona State University, P.O. Box 875906, Tempe, AZ 85287 5906, USA ^bLogistics Management Agency, United States Air Force, USA

Received 1 August 1999; received in revised form 1 September 2000; accepted 1 October 2000

Abstract

Logistics support is a key element of aircraft transportation systems. This paper is concerned with the impact of aircraft spares provisioning decisions on the availability of aircraft. Spares provisioning in this context is complicated by the fact that spares may be shared across aircraft and that aircraft may have redundant systems. In addition, decisions concerning aircraft spares support require a rapid response for safety reasons. Analytical models have proven to provide a quicker response time than corresponding simulation models. There is an existing analytical model that includes the effect of redundancy and spares, but the underlying assumption is that a large number of aircraft are being modeled. In many applications, predictions of the number of times an aircraft can fly each day and the number of aircraft that are ready at any time are applied to a small fleet of aircraft. This paper demonstrates the improvement in computational accuracy that is achieved by reflecting the impact of small numbers of aircraft on availability projections. The method is used to optimize spares provision with respect to a user specified availability goal. Although the case study for this work is a military combat aircraft application from the Gulf War, the method is applicable to any small system of vehicles or machines where components may be redundant, demand and repairs may be approximated as following an exponential distribution, and limited access to spare parts is the rule.

Scope and purpose

In many situations in exploration, mining, rescue, and defense, it is necessary to dispatch a small fleet of machines to a remote area to perform an important function. Since the location is difficult to access, resupply is often difficult or impossible. In anticipation of this situation, the machines often include redundant parts to allow for some component failures in the field that do not eliminate a machine from further use. Decisions have to be made, in advance, as to which backup components (spares) should be included with the machines

* Corresponding author. Tel.: + 1-480-965-3185; fax: + 1-480-965-8692. *E-mail address:* j.cochran@asu.edu (J.K. Cochran). at the remote site to be used when failures exceed the built-in redundancy level. This turns out to be an interesting modeling problem. Since the number of machines is relatively small, exact counting of machines and their components (both redundant and spares) must be done in order to make a decision. This leads to a combinatorial problem that cannot be solved for practical-sized problems with a reasonable computational effort by, say, discrete simulation. In this paper, we develop a methodology, based on finite queuing theory, which addresses this need. It is applied and validated on an availability situation similar to the deployment of aircraft to the Middle East during the Gulf War. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Logistics; Redundancy and spares; Queues; Simulation

1. Background

The United States Air Force currently uses the best analytical model available for the aircraft spares provisioning problem. It is named the Dynamic Multi-Echelon Technique for recoverable item control (Dyna-METRIC) developed by the RAND Corporation [1]. Dyna-METRIC is a flexible inventory modeling tool, but its underlying assumption is that a large number of entities are being modeled. The following is a brief discussion of the Dyna-METRIC model. The central theorem in the Dyna-METRIC model is Palm's theorem [2], also known as the infinite channel queuing assumption [3], which states that if demand for an item is given by a Poisson process with mean m per unit time, and if the repair time for each failed unit is independently and identically distributed according to any distribution with mean repair time T, then the steady-state probability distribution for the number of units in repair has a Poisson distribution with mean mT.

This is a powerful result since it does not depend on any assumption about underlying probability distributions, but when dealing with a small number of aircraft this theorem is violated because the demand distribution and the repair cycle are no longer independent. The problem is well documented and can occur with as many as eight aircraft [4].

The number and location of s spare parts in a supply system is

$$s = OH + DI - BO, \tag{1}$$

where OH is the number of spares on the shelf or on-hand, DI the number of items due-in from repair or resupply, and BO the number of items on backorder.

This is a balance equation where the order quantity is assumed to be one. Any change in one variable will result in a compensating change in another variable. For example, if an item demand occurs, the number of items due-in will increase by one and, if the current on-hand balance is greater than zero, then the on-hand balance will decrease by one. Otherwise, the number of backorders will increase by one.

Using the initial spares position of each item type *i*, the Dyna-METRIC model predicts the number of aircraft available and the number of sorties that can be flown each day. In order to make these predictions, the model must calculate the expected number of backorders $\text{EBO}_i(s_i)$ for each item type on the aircraft. This is done using

$$EBO_i(s_i) = x = \sum_{s_i+1}^{\infty} (x - s_i) Pr_i [DI = x],$$
(2)

530

where s_i is the stock level for each item type *i*, and $Pr_i[DI]$ the probability of DI items due in from repair or resupply for item type *i*.

An important common logistics measurement is "pipeline". Pipeline represents the number of units of an item in the repair or in the resupply chain. The average pipeline, μ , for the single base, full repair, no depot resupply case is the average demand, *m*, multiplied by the repair time, *T*, such that $\mu = mT$. As a result of Palm's Theorem, this average pipeline value becomes the mean of the Poisson distribution used to calculate the expected backorders. If we allow multiple bases, limited base repair, and depot repair and resupply, the average pipeline at base *j* becomes

$$\mu_j = m_j (r_j T_j + (1 - r_j) [O_j + \text{EBO}[s_0 | m_0 T_0] / m_0]),$$
(3)

where μ_j is the average pipeline at base j, m_j the average annual demand at base j, r_j the probability of repair at base j, T_j the average repair time at base j, O_j the average order and ship time from depot to base j, subscript j the base counter, and subscript o the depot counter.

For most single-unit aircraft combat assessments, this equation reduces to $\mu_j = (r_j T_j)m_j$ since depot repair and resupply are not available.

Aircraft availability (the ratio of aircraft available to aircraft fielded times 100) is an important measure of merit [5]. It can be calculated from the expected backorders as follows:

$$A = 100 \prod_{i=1}^{I} [1 - \text{EBO}_{i}(s_{i})/(NZ_{i})]^{Z_{i}}$$

subject to $EBO_i(s_i) \leq NZ_i$ for every item type *i*, (4)

where Z_i is the number of times the same item occurs on a single aircraft, and N the number of aircraft fielded.

This formula implies that an aircraft is available only when there are no failures in any of the Z_i items on an aircraft. The constraint prevents the number of backorders from exceeding the number of possible aircraft positions for each item type. The number of predicted aircraft flights/flying hours per day is simply the number of available aircraft each day times the maximum flight rate (flights/aircraft) per day which is capped at the total number of flights required each day.

This is a brief overview of the current model and a description of the current state of the modeling environment. Next, we introduce a method that addresses some of the shortcomings of the current method.

2. Research method

The following method is an effort to better model small systems of equipment where some of the assumptions made in Dyna-METRIC fail. However, the proposed method still maintains the advantages of the analytical model. The proposed method must overcome several challenges. These include the difficulties of a finite-calling population, the proper distribution of component back-orders, item redundancy, and the interdependence of component types. Previous researchers have dealt with the first two of these. We begin with these, then extend that redundancy work. Subsequently, we present a new iterative approach to handling interdependence.

Below is a failure distribution function from queuing theory that takes into account the finite calling nature of our situation [6]. Once all the original operating units and spares are consumed, no more arrivals (or failures) can occur. Also, as the number of operating components decreases, the arrival rate or failure rate naturally decreases as well. This is the probability distribution of failures, P(X):

$$P(X = x) = P(0)\frac{f^{x}R^{x}}{x!} \quad \text{when } x \leqslant s$$
(5)

and

$$P(0)\frac{f^{x}R^{s}}{x!}\frac{R!}{(R-x+s)!} \quad \text{when } s+1 \leq x \leq R+s,$$

where x is the number of component failures, f the ratio of failure rate to repair rate (λ/μ) , P(0) the normalizing constant so the distribution sums to one, s the number of spare components, and R the total number of each component being operated.

Notice that this result is only for one type of unit. In a situation where unit *i* is actually made up of r_i components and there are *M* such units in operation, the total number of operating components is $R = \sum_{i=1}^{M} r_i$.

To address the problem of component redundancy, Ref. [3] uses the hypergeometric distribution to properly distribute the total number of component failures, x, across all M operating units. The number of component failures, y_i , on each unit i is such that $\sum_{i=1}^{M} y_i = x$. The distribution of X is shown below:

hyp
$$(y_1, y_2, y_3, \dots, y_M | x) = \frac{\binom{r_1}{y_1}\binom{r_2}{y_2} \cdots \binom{r_M}{y_M}}{\binom{R}{x}},$$
 (6)

where x is the total number of component failures, r_i the number of components on each unit *i*, and y_i the number of component failures on unit *i* such that $\sum_{i=1}^{M} y_i = x$.

In this paper, redundancy can imply two things. Firstly, within a unit *i*, redundancy implies that not all r_i components must operate for the unit to function properly. Thus, assume that for unit *i* to operate properly, only q_i of the r_i components must be operating. Secondly, system redundancy implies that in a system of *M* units, not all *M* units need to operate for the system to function properly. Therefore, assume that *I* units of the *M* must be operating as a minimum for system success.

To illustrate these concepts more clearly, let us look at a simple radio example (see Fig. 1). This radio system is composed of transmitting and receiving units. Transmitter units for each radio are made up of three transmitter components ($r_1 = 3$). Receiver units are each made up of a single receiver component ($r_2 = 1$).

Based on the earlier discussions, the failure distribution of the units in this M radio system can be calculated. First, a failure distribution is needed for each unit type. Here we take the transmitter as one type (j = 1) and the receiver as a second type (j = 2). Since numbering of units is arbitrary, we assume the "first" fails. The failure distribution of unit type j, for the first 1, 2, ..., k units up and the

532



Fig. 1. M radio operating system.

 $k + 1, k + 2, \dots, M$ units down, is calculated using

$$U_{j}(k) = \sum_{x=0}^{\infty} P_{j}\{X=x\} \sum_{T} \text{hyp}(y_{1}, y_{2}, \dots, y_{M}|x), \quad I \leq k \leq M,$$
(7)

where j is the unit type number, T the complete set of component failure combinations that result in the first k units of type j up and the rest down, and I the minimum number of radios that must be available.

One advantage to this approach is that the set T of possible failure combinations, resulting from the first k units operating and the last M - k being down, limits the number of hypergeometric calculations that must be performed. Also, since these calculations are not dependent on the failure or repair distributions, they can be calculated once and used for all components.

The results for each unit type are then combined into a radio system failure distribution. For the radios, there are two different unit types, a transmitter and receiver type. The probability that all M radios are up and none are down, U(M), is given by

$$U(M) = U_1(M)U_2(M),$$
(8)

which is simply the probability of all the transmitter units being up times the probability of all the receiver units being up.

Similarly, U(M-1) is the probability that the first M-1 radios are up and radio M is down.

$$U(M-1) = U_1(M-1)[U_2(M-1) + U_2(M)] + U_1(M)U_2(M-1).$$
(9)

This equation results from unit combinations (transmitter and receiver) that create the first M - 1 radios up and the *M*th radio down. For example, the first term in Eq. (9) is created by the cases where either both unit types have the first M - 1 units up or the transmitter unit type has the first M - 1 up and the receiver unit type has all *M* up. Both of these combinations result in the first M - 1 radios up and the *M*th radio down. The equation for the first M - 2 radios up and the M - 1 and M radios down is more complex:

$$U(M-2) = U_1(M-2)[U_2(M-2) + 2U_2(M-1) + U_2(M)] + 2U_1(M-1)[U_2(M-2) + U_2(M-1)] + U_1(M)U_2(M-2).$$
(10)

The logic for building this equation is the same as explained previously, that is, enumerate all combinations that produce the first M - 2 radios up and the M - 1 and Mth radios down. For example, the second term $(U_1(M-2) \times 2U_2(M-1))$ is a result of the fact that with two radios down for transmitter units, there are two ways to select the radio that is also down for the receiver unit. This combinatorial equation building process continues until the minimum number of acceptable operating radios, I, is reached.

Based on the calculations for operating radio probabilities, an overall radio system availability can be calculated. Since U(k) is the probability that the first k radios are operating while the remaining M - k are down, the availability, A(k), of at least k radios operating can be found by computing the number of ways in which k up radios can be selected from the M radios. This results in the following equation:

$$A(k) = 100 \left[U(M) + \binom{M}{M-1} U(M-1) + \dots + \binom{M}{k} U(k) \right], \quad I \le k \le M.$$

$$\tag{11}$$

This is the only model, to date, that incorporates finite source arrivals, component redundancy, and multiple components into one availability model.

The preceding approach handles finite-population demand (i.e. the failure rate decreases as failures increase) within a component-type, but fails to consider the impact the failures that each component type have on the other component-type failures. Therefore, some type of adjustment must be made to account for this inter-component interaction. We propose an iterative approach. An adjustment of the other component-type's failure rates based on the highest failing component-type's expected availability will reflect this interaction. The expected number of operating units, $E[O_j]$, is calculated for each component-type j using

$$E[O_j] = \left[\sum_{k=0}^{M} \binom{M}{M-k} U_j (M-k) (M-k)\right].$$
(12)

Then the component type with the lowest number of expected operating units would be used to adjust the demand rates of the remaining component types. The adjustment factor is the availability of this least available component type. This availability is then multiplied by the arrival rate of

534



Fig. 2. Flow of model calculations.

all the other component types to obtain a new utilization rate. This equation is shown below:

$$f_{j\,\text{new}} = f_{j\,\text{old}} E[O_j]/M. \tag{13}$$

This new utilization rate, $f_{j \text{ new}}$, is used to recompute all the other component-type's $U_j(k)$. Then the system availability is recalculated, checking to see if the increase in A(M) (i.e. $A_{\text{Current}}(M) - A_{\text{Last}}(M)$) is more than 0.001 (we find 0.1% to work well). If it is, then the second lowest $E[O_j]$ value is selected and new $f_{j \text{ new}}$ values are computed for all but the lowest two $E[O_j]$ values. Then new availability's are computed. When the change in A(M) is less than 0.001, stop the process and report the availability results.

This iterative procedure requires only a small fraction of the computational time necessary by other techniques (such as simulation) to consider all of the combinations of failure interactions by focusing on the most likely to fail components. If we had a system with a single class of unit, and no spares, we could exploit the k-out-of-n system structure [7], but our system is more complicated than that. Our iterative procedure dramatically reduces computational effort while including both redundancy and spares. Fig. 2 and the algorithm given below summarize the method.

Algorithm.

Step 1: Calculate the component failure distributions, P(X), and the hypergeometric distribution of failures Hyp(y/x) for each component (Eqs. (5), (6)).

Step 2: Calculate the probability that the first k components of each are up and the k + 1, k + 2, ..., M are down, $U_j(k)$ (Eq. (7)).

Step 3: Calculate the probability that the first K units of the system are up and the K + 1, K + 2, ..., M units are down, U(k) (Eqs. (8)-(10)).

Step 4: If first pass or $A_{\text{current}}(M) - A_{\text{last}}(M) > 0.001$, go to step 5. Otherwise, go to step 8.

Step 5: Calculate the expected number of operating components of each component type, $E[O_j]$ (Eq. (12)).

Step 6: Select the smallest $E[O_j]$ value in step 5 and calculate the new utilization rates, $f_{j \text{ new}}$, for all the other component-types (Eq. (13)).

Step 7: Go to step 1 and repeat the steps for all components except the components whose $E[O_j]$ have already been computed.

Step 8: Calculate the availability, A(k), for the system based on the desired number of operating pieces in the system (Eq. (11)).

We now have a method to overcome all the shortcomings of current logistics models for spares provisioning. The next section presents computational and validation results based on a Gulf War case study.

3. Case study

Validation is always an important aspect of model development. In order to build a comparison database, a simulation model was developed using the SLAM II language [8] to create simulated data. The simulation model in [9] was used as a baseline for creating this simulated data. It had already been validated against operational data, so it provides an excellent starting point.

In addition, Headquarters Air Combat Command (HQ ACC) provided expert advice on data, modeling and validation issues. Aircraft in ACC are basically divided into two groups, fighters and "heavies". Fighters include aircraft such as the F-15, F-16, A-10 and F-111. Heavies include aircraft such as the B-52, B-1, E-3, and B-2. The test parameters were established through work with HQ ACC and are documented in [10]. Two aircraft from each group based on the recommendations of HQ ACC have been selected as typical of a deployed Gulf War unit. These aircraft are the B-52, E-3B, F-15, and F-16. For each of these aircraft types, three aircraft sizes were selected. Stock positions were set at three different levels: zero stock, fully authorized level, and the average fill percentage level based on March 1995 fill rates. Flight profiles were also provided by HQ ACC.

The results presented in Table 1 and Table 2 are a typical sample of the 120 different cases we analyzed. In only three cases of 120 were the proposed model results outside of the 95% confidence intervals for the simulation model, and the largest 95% simulation confidence interval based on availability was less than 4% wide. In addition, hypothesis of means tests were performed on the two different redundancy case output differences and they both had high-power values. This provides strong evidence that the new model is highly accurate. Furthermore, the run times for the analytical model required CPU seconds to execute, whereas the simulation model requires CPU hours to accomplish the same task.

4. An optimizing tool

Another capability missing from the current models is the ability to quantify any additional capability (or shortage) that might be available if the current availability goal is exceeded (or not met). For example, Dyna-METRIC will tell the user that he can fly his requested flying hours and maintain a 85% aircraft availability. But what if the availability goal is only 80%? Dyna-METRIC cannot tell you how many additional flying hours can be flown and yet still meet the availability goal. By adding a feedback mechanism to the proposed model and providing a user-input availability goal (Y), the proposed model can report the maximum operating hours per day per system, H, that can be achieved and still maintain the user availability goal (Y). This can be stated as

Maximize $H = OHC \times CPD$ Subject to $A(M-1) \ge Y$ for $M \ge 3$, $A(M) \ge Y$ for M < 3

where H is the number of operating hours per system per day, Y the user-input availability goal, A(M) and A(M - 1) are defined in Eqs. (8), (9), OHC the operating hours per cycle, and CPD the cycles per day.

This problem is solved by adding the following step to the algorithm presented previously.

Step 9: If the availability of M - 1 or more systems is greater than the availability goal, Y (the availability of M systems is used for scenarios where M < 3), increment the daily operating hours per system by the user provided step size, G, and go to step 1. If the availability goal, Y, is not exceeded, then report the availability and operating hours for the last acceptable operating hour/availability combination. If no such combination exists, simply report this operating hour program and the expected system availabilities, A(k)'s.

Of course, if the availability goal could not be met by the current stock quantities, the stopping rule is changed to decrement operating hours to find the maximum operating hours in order to meet the availability goal.

The following example illustrates how this capability might be used. In Table 1, the B-52 three aircraft "no redundancy" availability of two or more aircraft was 90.01%. What if the availability goal was only 80%? How many additional flying hours per aircraft per day could someone achieve and still meet the availability goal? If this information is used in the optimization model, we discover that instead of 7.2 flying hours per aircraft, we can get 10.20 flying hours per aircraft while achieving a 81.44% availability. This process could also be used to study the relationship between different stock levels and flying hours and their impact on aircraft availability using the optimization model.

Fig. 3 shows a complete response surface for the relationship between different stock levels, flying hours, and aircraft availability for the B-52 three aircraft case discussed above. Notice that it is unimodal (a property we observed in all cases), so that the optimization model yields the global solution (see Table 2).

Here we have a complete view of the general relationship between stock levels, flying hours, and availability which we could not create with current modeling capability. This graph was generated by running the model with 1% availability goals and a step size of 0.5 flying hours. The stock levels

	# Of A/C	Research model	Model error	Simulation 95% C.I.		
				Lower	Mean	Upper
E3B	5	19.40	0.06	18.14	19.34	20.54
		37.50	-0.06	36.49	37.56	38.62
	3	48.28	0.07	46.91	48.21	49.51
		39.71	0.06	38.77	39.65	40.53
	1	82.38	-0.43	81.84	82.81	83.79
F15	5	24.33	0.99	22.35	23.34	24.33
		39.67	0.22	38.66	39.45	40.24
	3	53.90	1.18	50.99	52.72	54.45
		36.94	-1.09	36.52	38.03	39.53
	1	86.89	-0.23	85.70	87.12	88.54
B52	5	25.78	- 0.33	25.34	26.11	26.89
		37.97	-0.73	38.20	38.70	39.20
	3	54.34	0.07	53.29	54.27	55.25
		35.67	-0.60	35.65	36.27	36.89
	1	82.38	-0.09	81.45	82.47	83.49
F16	5	87.97	-0.42	87.49	88.39	89.29
		11.41	0.48	10.11	10.93	11.75
	3	94.76	-0.46	94.50	95.22	95.93
		5.142	0.459	3.997	4.683	5.368
	1	97.73	0.27	96.89	97.46	98.04

Table 1 No redundancy comparison of results — full stock case



Fig. 3. Relationship of flying hours and stock level to availability.

	# Of A/C	Research model	Model error	Simulation 95% C.I.		
				Lower	Mean	Upper
E3B	5	17.67	0.07	16.81	17.60	18.38
		36.55	0.34	35.49	36.21	36.93
	3	45.47	0.26	43.58	45.21	46.85
		40.93	0.44	39.19	40.49	41.79
	1	77.63	-0.04	76.3	77.67	79.04
F15	5	19.36	0.54	17.96	18.82	19.67
		37.58	0.1	36.67	37.48	38.28
	3	46.46	0.71	43.96	45.75	47.54
		40.55	-0.13	39.21	40.68	42.16
	1	78.46	1.01	76.06	77.45	78.94
B52	5	16.31	0.46	14.96	15.85	16.74
		35.19	0.26	33.97	34.93	35.89
	3	56.18	-0.2	55.05	56.38	57.7
		34.61	-0.35	33.94	34.96	35.98
	1	82.72	- 0.3	82.03	83.02	84.01
F16	5	76.52	0.31	74.64	76.21	77.77
		20.88	0.08	19.46	20.8	22.15
	3	85.56	-0.48	85.05	86.04	87.04
		13.69	0.42	12.33	13.27	14.21
	1	95.49	0	94.69	95.49	96.29

 Table 2

 Redundancy case comparison of results — partial stock case

were decremented from 100% by randomly selecting the missing items for each additional loss in available stock until zero was reached. This graph allows the general assessment of any combination of flying hour and stock level decisions. However, it is much more computational demanding than a single search using the optimization model.

5. Summary

Logistics support has always been a key element of combat effectiveness as well as an important element in commercial aircraft systems. The Air Force is interested in developing tools and methods to assess the impact of logistics on combat capability. The current method used by the Air Force to assess aircraft spares support was developed for use with a large number of aircraft. With the high cost of new aircraft, fewer aircraft are fielded and, therefore, some of the critical assumptions made in the current model are unsatisfactory. This includes the assumption of an infinite calling population of demands and the distribution of backorders by sampling with replacement. This research develops an original method to assess the system availability of a small

number of vehicles or machines which overcomes the flaws of the existing model. The new method also allows the user to include component redundancy and component spares in the system.

In order to test the research method, a FORTRAN program was written to implement the approach. Using actual US Air Force failure, repair, and scenario data for the F-16, F-15, E-3B, and B-52H, the results were compared to simulation results. The research model was tested under two different cases: component redundancy and no component redundancy. In both cases, the research model performed very closely to the simulated data. This provides strong evidence that the research method performs well under a wide range of operating conditions. The research model has an efficiency advantage over the simulation model which include speed of processing (performance improvement of several orders of magnitude) and, thus, the ability to handle more components in the system for a fixed response time.

Finally, an additional capability was added to the research method that does not exist in the current model or in the simulation model. An optimization technique is provided that allows the user to optimize operating hours given a system availability goal or target. This new ability allows the user to estimate the additional capability available (i.e. additional operating hours), or shortage, if he is currently not matching his availability goal.

References

- [1] Isaacson K, Boren P, Tsai C, Pyles R. Dyna-METRIC Version 4: Modeling Worldwide Logistics Support of Aircraft Components. The RAND Corp., R-3389-AF, 1988.
- [2] Palm C. Analysis of the Erlang Traffic Formulae for Busy-Signal Arrangements. Ericsson Technics 1938;4:39–58.
- [3] Sherbrooke C. Optimal inventory modeling of systems: multi-echelon techniques. New York: Wiley, 1992. p. 274.
- [4] Faughaber K. DMAS small PAA utility. Washington, DC: HQ USAF/LGSS, 1993.
- [5] Pyles R, Tripp R. Dyna-metric: a valuable management tool. Air Force Journal of Logistics 1983;18-21:24.
- [6] Kaplan A. Incorporating redundancy considerations into stockage policy models. Naval Research Logistics 1989;36:625-38.
- [7] Rushdi AM. Utilization of symmetric switching functions in the computation of K-out-of-N system reliability. Microelectronics and Reliability 1986;26(5):973-87.
- [8] Pritsker A. Introduction to simulation and SLAM II. New York: Wiley, 1986.
- [9] Lewis T. Unit level WRSK assessment and sortie generation simulation model. MS thesis, Air Force Institute of Technology, Wright Patterson AFB, OH, 1987.
- [10] Miyares CP. Validation of aircraft combat capability model memorandum. HQ ACC/LGSW, Langley AFB, VA., 24 Mar 95, 1995.

Jeffery K. Cochran is Professor of Industrial Engineering at Arizona State University. Dr. Cochran received his Ph.D. in the area of Operations Research from Purdue University in 1984. His industrial experience includes positions at Battelle Northwest Laboratory, Los Alamos National Laboratory, and NASA's Laboratory for Applications of Remote Sensing. Professor Cochran has directed research projects for the National Science Foundation's I/UCRC for Manufacturing and Automation, the National Institute of Standards and Technology, and the Institute for Manufacturing and Automation Research. His primary research interests are optimization of stochastic models in high technology production systems, next-generation discrete event simulation tools, and transport system and supply chain logistics. Dr. Cochran is the author of over ninety scholarly publications, and is a Professional Engineer.

Theodore P. Lewis is a Major in the United States Air Force and a Technical Director of the USAF Logistics Management Agency. Dr Lewis received his Ph.D. in Industrial Engineering from Arizona State University in 1995.